The Banach-Tarski Paradox

Esben Bistrup Halvorsen   Lars Hougaard Hansen

October 8, 2002

Contents

Preface

This bachelor thesis concentrates on proving the famous Banach–Tarski paradox and Tarski's theorem, and is a result of 4 months of work under the adept guidance of our adviser, Mikael Rørdam. We intend that students at the final stages of an undergraduate study in mathematics can read and understand the thesis. Our presentation presupposes acquaintance with topology and elementary group theory. Furthermore, some knowledge of measure theory could ease the conception of Tarski's theorem.

The text is divided into three chapters. Chapter 1 introduces the concepts of paradoxicality and free groups, which are necessary in our proofs of the Banach–Tarski paradox and Tarski's theorem. Furthermore, Chapter 1 includes a section in which we investigate the free product of groups, which, although it is not being used directly, is essentially what lies behind some of the proofs leading to the Banach–Tarski paradox. Chapter 2 focuses on the Banach–Tarski paradox, both in its original form and in an even stronger version. Chapter 3 extends some of the ideas from the first chapter and proves Tarski's theorem.

The basis for our work is the excellent book, *The Banach–Tarski Paradox*, by Stan Wagon. Since many proofs in this book have either been omitted or are very short, we have deepened Stan Wagon's presentation. Furthermore, Stan Wagon only superficially discusses the theory of free groups and we disagree with his idea of this concept. We therefore had to reformulate some of the definitions and theorems and, in some cases, had to construct new proofs.

Finally, we express our gratitude towards Mikael Rørdam for generously sharing his sparse time with us and for his commitment to the process and to David Jeffrey Breuer for editing the text.

Esben Bistrup Halvorsen
Lars Hougaard Hansen

Introduction

Phenomena that seem counterintuitive at first sight appear in most aspects of modern mathematics. Nearly all paradoxes in mathematics are connected with the inconceivable concept of *infinity*, which throughout history has challenged the minds of human beings. The idea, proposed by Cantor in the late 19th century, that the difference in the cardinality of various infinite sets indicated that there were actually several types of infinity, was originally met with great scepticism. However, as time went by, mathematicians adjusted to this preposterous thought and learned to accept (and maybe even love) the unfathomable features of the infinite.

A source of manifold paradoxes is the intuitively acceptable *axiom of choice*, which was originally formulated in the early 20th century by Zermelo.

The Axiom of Choice *For every nonempty set M, there exists a mapping*

$$u\colon \mathcal{P}(M)\backslash\{\varnothing\} \to M$$

such that

$$\forall A \in \mathcal{P}(M)\backslash\{\varnothing\} : u(A) \in A.$$

In this paper, theorems proved using the axiom of choice will be followed by (AC). The axiom of choice plays an important role in proving some of the most fundamental results in modern mathematics, such as Hahn–Banach's theorem and Tychonoff's theorem, and in 1938 Gödel proved it to be consistent with the Zermelo–Fraenkel axioms. However, the apparently harmless axiom leads to paradoxes so absurd that one is easily tempted to reject it, and after Cohen's proof from 1963 that also the *negated* axiom of choice is consistent with the Zermelo–Fraenkel axioms, the axiom has engendered a fundamental discussion in the world of mathematics. One of the astounding consequences of the axiom of choice is the Banach–Tarski paradox:

Theorem (The Banach–Tarski Paradox)(AC) *There exists a partitioning of the unit ball B into pieces $A_1, \dots, A_n$, $B_1, \dots, B_m \subseteq \mathbb{R}^3$ and isometries $\phi_1, \dots, \phi_n$, $\psi_1, \dots, \psi_m$ of $\mathbb{R}^3$, such that*

$$B = \bigcup_{i=1}^{n} \phi_i(A_i) = \bigcup_{j=1}^{m} \psi_j(B_j).$$

The paradoxical nature of the Banach–Tarski paradox is more apparent in its less formal formulation: It is possible to cut an orange into a finite number of pieces that can be rearranged to form two oranges with the same size as the original one!

The Banach–Tarski paradox is formulated and proved here with the aid of the theory of *paradoxicality*. The development of paradoxicality began with the formalization of measure theory in the early 20th century. One of the first instances of the use of a paradoxical decomposition is Vitali's classical example from 1905 of a non-Lebesgue measurable set. In 1915, Hausdorff proved the non-existence of another type of measure by constructing a truly surprising paradox, and this inspired some important work in the 1920s, such as the Banach–Tarski paradox from 1924. The connection between paradoxicality and measure theory is emphasized by Tarski's theorem, which states that a necessary and sufficient condition for the existence of a paradoxical decomposition of a set is the absence of a finitely additive and invariant measure, normalizing the set.

Chapter 1

Paradoxicality and Free Groups

1.1 Paradoxical Actions

In order to prove the Banach–Tarski paradox, we need to formulate it in terms of paradoxicality, which again is defined in the context of group actions. So recall that a group G is said to act on a set X if there exists a mapping $G \times X \to X$, denoted $(g, x) \mapsto g.x$, such that for all $g, h \in G$ and $x \in X$,

$$1.x = x \qquad (gh).x = g.(h.x),$$

where 1 is the neutral element of G. We shall by the orbit of x, where $x \in X$, be referring to the set

$$G.x = \{g.x \mid g \in G\}.$$

Definition 1.1 *Let G be a group acting on a set X and let $E \subseteq X$ be a subset of X. E is G-*paradoxical *(or paradoxical with respect to G) if there exist pairwise disjoint subsets $A_1, \ldots, A_n, B_1, \ldots, B_m$ of E and elements $g_1, \ldots, g_n, h_1, \ldots, h_m \in G$ such that*

$$E = \bigcup_{i=1}^{n} g_i.A_i = \bigcup_{j=1}^{m} h_j.B_j.$$

A group G is said to be *paradoxical* if it is G-paradoxical, where G acts on itself by left multiplication. Even though the concept of paradoxicality may seem abstract, a wide quantity of well-known examples satisfy the conditions in Definition 1.1.

Example 1.2 *The even numbers E are paradoxical with respect to the group of bijections on $\mathbb{N}$. This can be realized by dividing E into classes, E_1, E_2, with respect to congruence modulo 4. If we let O be the odd numbers, then, as $|E_1| = |E|$ and $|E_2 \cup O| = |O|$, there exist bijections, $f_1 \colon E_1 \to E$ and $f_2 \colon E_2 \cup O \to O$. By defining $f \colon \mathbb{N} \to \mathbb{N}$ by*

$$f(n) = \begin{cases} f_1(n) & \text{for } n \in E_1 \\ f_2(n) & \text{for } n \in E_2 \cup O \end{cases}$$

we see that f is a bijection on $\mathbb{N}$ with $f(E_1) = E$. Equivalently, one can define a bijection g on $\mathbb{N}$ such that $g(E_2) = E$. Hence,

$$E = f(E_1) = g(E_2).$$

Note that the axiom of choice is not involved in the above example, so the concept of paradoxicality is not solely related to this. However, the most interesting cases of paradoxicality are connnected to the axiom of choice.

With the aid of paradoxicality, we can now reformulate the Banach–Tarski paradox:

Theorem 1.3 (The Banach–Tarski Paradox)(AC) *The unit ball B in $\mathbb{R}^3$ is SO_3-paradoxical.*[1]

The paradoxicality of a group can, in a sense, be transferred to the set it acts on. This is contained in the following theorem.

Theorem 1.4 (AC) *Let G be a group that acts on the set X without nontrivial fixed points (that is, $g.x = x \Leftrightarrow g = 1$). Then G is paradoxical if, and only if, X is G-paradoxical.*

Proof: Suppose $A_1, \dots, A_n, B_1, \dots, B_m$ are disjoint subsets of G and $G = \bigcup_{i=1}^n g_i.A_i = \bigcup_{j=1}^m h_j.B_j$, where $g_1, \dots, g_n, h_1, \dots, h_m \in G$. Letting u be the choice function on X, we define $M = \{u(G.x) \mid x \in X\}$. If $g.x = h.y$ for some $g, h \in G$ and $x, y \in M$, then, since y and $x = (g^{-1}h).y$ belong to the same orbit, $x = y$. Furthermore, as the action of G on X has no nontrivial fixed points, $g.x = h.x$ implies that $g = h$. Hence, $g.M \cap h.M \neq \varnothing$ leads to $g = h$, that is, the sets $g.M$, $g \in G$, are pairwise disjoint. To see that X is covered by the sets, suppose $x_0 \in X$. Then there exists $g \in G$ such that $g.x_0 = u(G.x_0)$, and we find that $x_0 = g^{-1}.u(G.x_0) \in g^{-1}.M$, so the sets $g.M$, $g \in G$ partition X. Now, let

$$A_i' = \bigcup_{g \in A_i} g.M = A_i.M \quad \text{and} \quad B_j' = \bigcup_{g \in B_j} g.M = B_j.M$$

for $i = 1, \dots, n$ and $j = 1, \dots, m$. Then, since $A_1, \dots, A_n, B_1, \dots, B_m$ are pairwise disjoint, so are $A_1', \dots, A_n', B_1', \dots, B_m'$. Now we have

$$\begin{aligned} \bigcup_{i=1}^n g_i.A_i' &= \bigcup_{i=1}^n g_i.(A_i.M) \\ &= (\bigcup_{i=1}^n g_i A_i).M \\ &= G.M \\ &= X \end{aligned}$$

[1] SO_n denotes the group of isometries in $\mathbb{R}^n$.

Similarly, we obtain $\bigcup_{j=1}^m h_j.B'_j = X$. Hence, X is G-paradoxical.

Conversely, suppose X is G-paradoxical by $X = \bigcup_{i=1}^n g_i.A_i = \bigcup_{j=1}^m h_i.B_j$ and consider, for some fixed $x_0 \in X$, the orbit $G.x_0$. Let

$$T_i = \{g \in G \mid g.x_0 \in A_i\} \quad \text{and} \quad S_j = \{g \in G \mid g.x_0 \in B_j\}$$

for $i = 1, \ldots, n$ and $j = 1, \ldots, m$. Note that the sets are pairwise disjoint, since they are the counterimages under the mapping $g \mapsto g.x_0$ of the disjoint A_i's and B_j's. Since $G.x_0 = \bigcup_{i=1}^n (g_i.A_i) \cap G.x_0$, we find for any element $g \in G$, that $g.x_0 = g_i.y$ for some i and some $y \in A_i$. This implies that $g_i^{-1}g \in T_i$, and we conclude that $g \in g_iT_i$. Hence, $G = \bigcup_{i=1}^n g_iT_i$. Similarly, $G = \bigcup_{j=1}^m h_jS_j$, so G is paradoxical.

□

Note that the proof of the last implication does not require the assumption that the action of G has no nontrivial fixed points, so we have actually proved that, if there exists a set X such that X is G-paradoxical, then G is itself paradoxical.

The action of a group on itself obviously has no nontrivial fixed points. Consequently, the action of a subgroup on the entire group has no nontrivial fixed points. Theorem 1.4 now gives us:

Corollary 1.5 (AC) *A group with a paradoxical subgroup is paradoxical.*

An important tool in our proof of the Banach–Tarski paradox is Theorem 1.4, as we shall attempt to find a paradoxical subgroup of SO_3 that acts on the unit ball B in $\mathbb{R}^3$.

1.2 Free Groups

The group we wish to use in creating the paradoxicality of B is a *free group*. As the reader is not expected to be familiar with free groups, we introduce this important concept.

Definition 1.6 *A group G is said to be free on a subset $S \subseteq G\backslash\{1\}$ if any element $g \in G\backslash\{1\}$ has a unique representation*

$$g = y_1 y_2 \cdots y_k, \tag{1.1}$$

with $y_i \in \langle x_i \rangle \backslash \{1\}$, $x_i \in S$ for $i = 1, \ldots, k$ and $x_j \neq x_{j+1}$ for $j = 1, \ldots, k-1$.

A group G is simply said to be free if it is free on a subset $S \subseteq G\backslash\{1\}$. In this case, S is said to freely generate G. We shall refer to the representation (1.1) as a word, and $y_1, \ldots, y_k$ as the letters of the word. A set S as in Definition 1.6 is called a basis for G. It can be proved that all bases are equipotent; we can therefore define the rank of a free group as the cardinality of one of its bases. Obviously, if two free groups G_1, G_2 have bases B_1, B_2 that are "isomorphic" in the sense that there exists a bijection $t\colon B_1 \to B_2$ with $|t(x)| = |x|$ for all $x \in B_1$, then G_1 and G_2 are isomorphic by $x_1^{n_1} \cdots x_k^{n_k} \mapsto t(x_1)^{n_1} \cdots t(x_k)^{n_k}$.

Example 1.7 *Suppose $S = \{\sigma, \tau\}$ is a basis for the free group G of rank 2. Furthermore, suppose $|\sigma| = |\tau| = 2$. Then the elements of G can be regarded as words in σ and τ. According to the first and last letter, there will be four kinds of words in G. Any product of two elements in G will produce either the identity 1 or an element written as in (1.1) after successive removal of identical adjacent letters. The identity is also denoted the empty word.*

Example 1.8 *The fundamental group of any subset of the plane shaped like an 8 is a free group with a basis consisting of a loop traversing the upper circle and a loop traversing the lower circle of the 8. Consequently, the group has rank 2 and generators of infinite order.*

One of the simple examples of a paradoxical group is the free group of rank 2 with generators of infinite order. It is this group we need in order to create the paradoxicality of the unit ball in $\mathbb{R}^3$.

Theorem 1.9 *Suppose F is a free group of rank 2, and let $S = \{\rho, \pi\}$ be a basis for F. If $|\rho| = |\pi| = \infty$, then F is paradoxical.*

Proof: Let, for $\lambda \in \{\rho, \rho^{-1}, \pi, \pi^{-1}\}$, $W(\lambda)$ denote the set of words beginning on the left with the letter λ^n, for some $n \in \mathbb{N}$. Then $W(\rho)$, $W(\rho^{-1})$, $W(\pi)$ and $W(\pi^{-1})$ are disjoint subsets of F and

$$F = W(\rho) \cup \rho W(\rho^{-1}) = W(\pi) \cup \pi W(\pi^{-1}),$$

since, for example, all $\sigma \in F \backslash W(\pi)$ satisfy the criterion that $\pi^{-1}\sigma \in W(\pi^{-1})$.

□

1.3 Free Product of Groups

From any two groups we can construct a free group called the free product. Though the free product is not used directly here, it is essentially what lies behind the forthcoming proof of the Hausdorff paradox, and we therefore introduce this concept to provide greater insight.

Definition 1.10 *Suppose G_1 and G_2 are groups. A* free product *of G_1 and G_2 is a group F with homomorphisms $\phi_1\colon G_1 \to F$ and $\phi_2\colon G_2 \to F$, satisfying the condition that for any group H and homomorphisms $\psi_1\colon G_1 \to H$ and $\psi_2\colon G_2 \to$*

H, there exists one, and only one, homomorphism $\xi\colon F \to H$, such that the diagram below commutes.

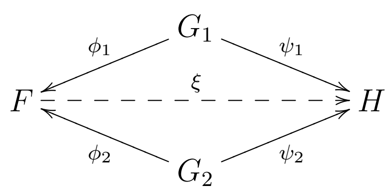

(1.2)

The definition does not fulfil our ambition of *the* free product, since neither the existence nor the uniqueness is apparent. The following propositions do the job.

Proposition 1.11 (uniqueness) *Suppose F and K are free products of G_1 and G_2. Then F and K are isomorphic.*

Proof: Let ϕ_1 and ϕ_2 witness that F is a free product of G_1 and G_2 and let λ_1 and λ_2 witness that K is a free product of G_1 and G_2. The hypothesis postulates that there exist homomorphisms, ξ and κ, such that the following diagrams commute.

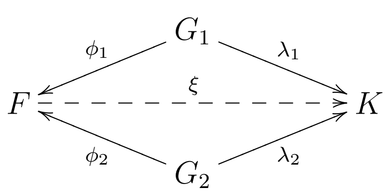

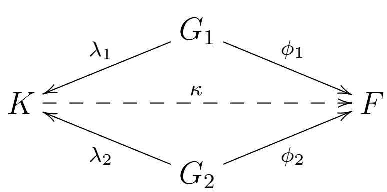

This shows that $\kappa \circ \xi \circ \phi_i = \kappa \circ \lambda_i = \phi_i$ for $i = 1, 2$, that is, the diagram below commutes.

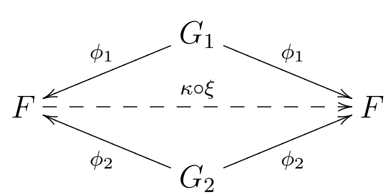

Replacing the homomorphism $\kappa \circ \xi$ with the identity, id_F, also makes the diagram commute, and since F is a free product, this implies that $\kappa \circ \xi = \mathrm{id}_F$. Similarly, $\xi \circ \kappa = \mathrm{id}_K$, and we conclude that F and K are isomorphic.

□

Proposition 1.12 (existence) *For any two groups G_1 and G_2, there exists a free product.*

Proof: Let F be the set of finite sequences $(c_1, \ldots, c_n)$, where the c_i's belong to $G_1 \backslash \{1\}$ and $G_2 \backslash \{1\}$ alternately, and where $n = 0$ gives the empty sequence, $()$. We inductively define a composition in F by $()() = ()$, $()(c_1, \ldots, c_n) =$

$(c_1, \ldots, c_n)() = (c_1, \ldots, c_n)$ and for $a = (c_1, \ldots c_n)$ and $b = (d_1, \ldots, d_m)$, where $n, m \geq 1$,

$$ab = \begin{cases} (c_1, \ldots, c_n, d_1, \ldots, d_m) & \text{if } c_n \in G_1, d_1 \in G_2 \text{ or } c_n \in G_2, d_1 \in G_1 \\ (c_1, \ldots, c_n d_1, \ldots, d_m) & \text{if } c_n, d_1 \in G_i,\ i = 1, 2, \text{ and } c_n d_1 \neq 1 \\ (c_1, \ldots, c_{n-1})(d_2, \ldots, d_m) & \text{if } c_n, d_1 \in G_i,\ i = 1, 2, \text{ and } c_n d_1 = 1 \end{cases}$$

It is easily verified that F with the given composition is a group. Now, let $\phi_1 \colon G_1 \to F$ and $\phi_2 \colon G_2 \to F$ be given by[2] $\phi_i(1) = ()$ and $\phi_i(g) = (g)$ for $g \in G_i \backslash \{1\}$. If H is a group and $\psi_i \colon G_i \to H$ are homomorphisms, then $\xi \colon F \to H$, given by $\xi(()) = 1$ and

$$\xi((c_1, \ldots, c_n)) = \begin{cases} \psi_1(c_1)\psi_2(c_2)\cdots\psi_1(c_n) & \text{if } c_1 \in G_1 \text{ and } c_n \in G_1 \\ \psi_1(c_1)\psi_2(c_2)\cdots\psi_2(c_n) & \text{if } c_1 \in G_1 \text{ and } c_n \in G_2 \\ \psi_2(c_1)\psi_2(c_2)\cdots\psi_1(c_n) & \text{if } c_1 \in G_2 \text{ and } c_n \in G_1 \\ \psi_2(c_1)\psi_2(c_2)\cdots\psi_2(c_n) & \text{if } c_1 \in G_2 \text{ and } c_n \in G_2 \end{cases}$$

is a homomorphism that makes diagram (1.2) commute. If ξ' is any homomorphism making the diagram commute, then

$$\begin{aligned} \xi'((c_1, \ldots, c_n)) &= \xi'(\phi_1(c_1)\phi_2(c_2)\cdots\phi_1(c_n)) \\ &= \xi'(\phi_1(c_1))\xi'(\phi_2(c_2))\cdots\xi'(\phi_1(c_n)) \\ &= \psi_1(c_1)\psi_2(c_2)\cdots\psi_1(c_n) \\ &= \xi((c_1, \ldots, c_n)) \end{aligned}$$

if $c_1, c_n \in G_1$. Similarly, $\xi'((c_1, \ldots, c_n)) = \xi((c_1, \ldots, c_n))$ in the other cases of c_1 and c_n, so $\xi' = \xi$, that is, ξ is the only homomorphism making the diagram commute. Hence, F is a free product of G_1 and G_2.

□

The free product of G_1 and G_2 is denoted $G_1 * G_2$. The construction of the free product in the above proof shows that $G_1 * G_2$ is a free group, as we can imagine the sequences as being words. The group $\mathbb{Z} * \mathbb{Z}$ is therefore free and has a basis consisting of two elements of infinite order. Consequently, $\mathbb{Z} * \mathbb{Z}$ is paradoxical according to Theorem 1.9.

[2]to ease notation, we use the same symbols for the neutral elements in G_1 and G_2.

Chapter 2

Equidecomposability and the Banach–Tarski Paradox

With the theory of free groups and paradoxicality well in place, we are now ready to sketch our strategy in proving the Banach–Tarski paradox. As mentioned in the previous chapter, our first goal is to find a paradoxical subgroup of SO_3. This leads to the Hausdorff paradox, which says that there exists a countable set $D \subseteq S^2$ such that $S^2 \backslash D$ is SO_3-paradoxical. Using the theory of equidecomposability, the Hausdorff paradox implies that S^2 is SO_3-paradoxical. A simple consequence of this is the Banach–Tarski paradox: B is SO_3-paradoxical.

2.1 The Hausdorff Paradox

Our way of finding a paradoxical subgroup of SO_3 could seem a little awkward. We will pretend that a certain subgroup F of SO_3 was sent to us from heaven. To show that F is paradoxical, we want to use Theorem 1.9 and therefore need to show that it is free. This will be done by considering an even larger subgroup T of SO_3 and showing that this is free. We will work our way backwards, and commence by showing that T is free.

Lemma 2.1 *There exist two rotations ϕ, ψ about axes through the origin in $\mathbb{R}^3$ that freely generate a subgroup T of SO_3.*

PROOF: Let

$$\psi = \begin{pmatrix} -\frac{1}{2} & -\frac{\sqrt{3}}{2} & 0 \\ \frac{\sqrt{3}}{2} & -\frac{1}{2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \phi = \begin{pmatrix} -\cos\theta & 0 & \sin\theta \\ 0 & -1 & 0 \\ \sin\theta & 0 & \cos\theta \end{pmatrix}$$

where θ is chosen so that $\cos\theta$ is transcendental. (This can be done since there are only countably infinite algebraic numbers, and the image of cosine is $[-1, 1]$.) Note that ψ and ϕ are both rotations about axes through the origin and that

$\psi^3 = \phi^2 = 1$. We want to show that, if $y = y_1 y_2 \cdots y_n$ and $z = z_1 z_2 \cdots z_m$ are words in $\{\phi, \psi\}$ and $y = z$, then $n = m$ and $y_i = z_i$ for $i = 1, 2, \ldots, n$.

Let us first realize, that it suffices to show that no word written as in (1.1) equals the identity 1. To see this, suppose that $y = z$, i.e. $yz^{-1} = 1$. Then $1 = y_1 y_2 \cdots y_n z_m^{-1} z_{m-1}^{-1} \cdots z_1^{-1}$. If either $n \neq m$ or $y_i \neq z_i$ for some i, then by successive cancellation of adjacent letters (beginning at $y_n z_1^{-1}$), we obtain a nonempty word that equals the identity.

To show that no word written as in (1.1) is equal to the identity, we divide the possible nonempty words into four types,

$$\alpha = \psi^{p_1}\phi\psi^{p_2}\phi\cdots\psi^{p_n}\phi, \qquad \beta = \phi\psi^{p_1}\phi\psi^{p_2}\cdots\phi\psi^{p_n},$$
$$\gamma = \phi\psi^{p_1}\phi\psi^{p_2}\cdots\psi^{p_n}\phi, \qquad \delta = \psi^{p_1}\phi\psi^{p_2}\phi\cdots\phi\psi^{p_n},$$

where $p_i \in \{1, 2\}$ for $i = 1, \ldots m$. We consider the word ϕ as of type γ and the word ψ as of type δ. To begin with, we consider words of type α. Notice that a word of type α can be written in the form

$$\alpha = \sigma_n \sigma_{n-1} \cdots \sigma_1,$$

where σ_i is either $\psi\phi$ or $\psi^2\phi$. Computing these two expressions gives

$$\sigma_i = \begin{pmatrix} \frac{1}{2}\cos\theta & \pm\frac{\sqrt{3}}{2} & -\frac{1}{2}\sin\theta \\ \mp\frac{\sqrt{3}}{2}\cos\theta & \frac{1}{2} & \pm\frac{\sqrt{3}}{2}\sin\theta \\ \sin\theta & 0 & \cos\theta \end{pmatrix}$$

for $i = 1, \ldots, n$.

By induction on n, we now want to show that, if $K = (0, 0, 1)$, then there exist rational polynomials P_{n-1}, Q_{n-1} and R_n, with $\deg P_{n-1} = \deg Q_{n-1} = n - 1$ and $\deg R_n = n$, and with leading coefficients $-\frac{1}{2}(\frac{3}{2})^{n-1}$, $\pm\frac{1}{2}(\frac{3}{2})^{n-1}$ and $(\frac{3}{2})^{n-1}$ respectively, such that

$$\sigma_n\sigma_{n-1}\cdots\sigma_1(K) = (\sin\theta P_{n-1}(\cos\theta), \sqrt{3}\sin\theta Q_{n-1}(\cos\theta), R_n(\cos\theta)).$$

For $n = 1$,

$$\alpha(K) = \sigma_1(K) = (-\tfrac{1}{2}\sin\theta, \pm\tfrac{\sqrt{3}}{2}\sin\theta, \cos\theta),$$

so we can choose $P_0(x) = -\frac{1}{2}$, $Q_0(x) = \pm\frac{1}{2}$ and $R_1(x) = x$. Assume next that the condition holds for $n = m - 1$. Then for $n = m$,

$$\begin{aligned} \alpha(K) &= \sigma_m\sigma_{m-1}\cdots\sigma_1(K) \\ &= \sigma_m(\sin\theta P_{m-2}(\cos\theta), \sqrt{3}\sin\theta Q_{m-2}(\cos\theta), R_{m-1}(\cos\theta)) \\ &= (\sin\theta P_{m-1}(\cos\theta), \sqrt{3}\sin\theta Q_{m-1}(\cos\theta), R_m(\cos\theta)), \end{aligned}$$

where we (inductively) have defined the polynomials

$$\begin{aligned} P_{m-1}(x) &= \tfrac{1}{2}xP_{m-2}(x) \pm \tfrac{3}{2}Q_{m-2}(x) - \tfrac{1}{2}R_{m-1}(x) \\ Q_{m-1}(x) &= \mp\tfrac{1}{2}xP_{m-2}(x) + \tfrac{1}{2}Q_{m-2}(x) \pm \tfrac{1}{2}R_{m-1} \\ R_m(x) &= (1 - x^2)P_{m-2}(x) + xR_{m-1}(x) \end{aligned}$$

with leading coefficients

$$\begin{array}{rll} P_{m-1}: & \frac{1}{2}(-\frac{1}{2}(\frac{3}{2})^{m-2}) - \frac{1}{2}(\frac{3}{2})^{m-2} & = -\frac{1}{2}(\frac{3}{2})^{m-1} \\ Q_{m-1}: & \mp\frac{1}{2}(-\frac{1}{2}(\frac{3}{2})^{m-2}) \pm \frac{1}{2}(\frac{3}{2})^{m-2} & = \pm\frac{1}{2}(\frac{3}{2})^{m-1} \\ R_m: & \frac{1}{2}(\frac{3}{2})^{m-2} + (\frac{3}{2})^{m-2} & = (\frac{3}{2})^{m-1} \end{array}$$

It follows further from this that the polynomials have the desired degrees, so the result is true for $n = m$, and hence by induction for all $n \in \mathbb{N}$. Now, since $\cos\theta$ is transcendental, we cannot have $\alpha(K) = K$, as this would require $R_n(\cos\theta) - 1 = 0$. We therefore conclude that $\alpha \neq 1$.

Next assume that a word of type β is equal to 1. Then the word $\phi\beta\phi = \phi 1 \phi = 1$ is of type α, which is a contradiction. So words of type β are also different from 1.

Assume now that $\delta = \psi^{p_1}\phi\psi^{p_2}\phi\cdots\phi\psi^{p_m} = 1$, and that m is the smallest positive integer for which this is true. Since $\psi \neq 1$, we have $m > 1$. If $p_1 = p_m$, then

$$1 = \psi^{-p_1}\delta\psi^{p_1} = \phi\psi^{p_2}\cdots\phi\psi^{p_1+p_m}$$

is a word of type β, since $\psi^{p_1+p_m} = \psi^2$ or $\psi^{p_1+p_m} = \psi^4 = \psi$; but we have already shown that this is not possible. If $p_1 \neq p_m$, then $p_1 + p_m = 3$. First consider $m > 3$. Then

$$1 = \phi\psi^{p_m}\delta\psi^{p_1}\phi = \psi^{p_2}\phi\cdots\phi\psi^{p_{m-1}},$$

which contradicts the minimality of m. On the other hand, we cannot have $m = 2$, because this would imply $1 = \psi^{p_2}\delta\psi^{p_1} = \phi$. The last possible case is $m = 3$, but this gives $1 = \phi\psi^{p_3}\delta\psi^{p_1}\phi = \psi^{p_2}$. In all cases, 1 is not a word of type δ.

Finally, $\gamma = \phi\psi^{p_1}\phi\psi^{p_2}\ldots\psi^{p_m}\phi = 1$ implies that $\phi\gamma\phi = \psi^{p_1}\phi\psi^{p_2}\ldots\psi^{p_m} = 1$ for $m \geq 1$, that is, 1 is a word of type δ. As $\phi \neq 1$ is obvious, we conclude that all words of type γ are different from 1.

Hence, no word is equal to 1, and based on the introductory remarks we conclude that ψ and ϕ freely generate a subgroup of SO_3.

□

Now that we have shown that T is a free subgroup of SO_3, we can reveal that our candidate for F is the set of words generated by $x_1 = \psi\phi\psi$ and $x_2 = \phi\psi\phi\psi\phi$. It will then be shown that almost the entire unit sphere S^2 is F-paradoxical and therefore is SO_3-paradoxical.

Theorem 2.2 (The Hausdorff Paradox)(AC) ***(AC)*** *There exists a countable subset D of the unit sphere S^2, such that $S^2 \backslash D$ is SO_3-paradoxical.*

Proof: Consider the two rotations $x_1 = \psi\phi\psi$ and $x_2 = \phi\psi\phi\psi\phi$. We claim that x_1 and x_2 freely generate a subgroup F of T. Notice that $x_1^n = \psi\phi\psi^2\cdots\psi^2\phi\psi$ and $x_2^n = \phi\psi\phi\psi^2\cdots\psi^2\phi\psi\phi$ are nontrivial words in $\{\phi, \psi\}$ for all $n \in \mathbb{N}$, so, since

T is free on $\{\phi, \psi\}$, x_1 and x_2 are of infinite order. Now, let $F = \{z_1 z_2 \cdots z_k \mid z_i \in \langle x_1 \rangle \cup \langle x_2 \rangle\}$. It is obvious that F is a subgroup of T. To show that F is free on $\{x_1, x_2\}$, as in the proof of Lemma 2.1, we only need to show that no word writen as in (1.1) is equal to 1. So, to get a contradiction, assume that

$$1 = y_1 y_2 \cdots y_k$$

is a word in $\{x_1, x_2\}$. Since any power of x_1 begins and ends with a ψ and any power of x_2 begins and ends with a ϕ, the element $y_1 y_2 \cdots y_k$ is a nonempty word in $\{\phi, \psi\}$, which contradicts T being free on $\{\phi, \psi\}$. We conclude that F is a free group of rank 2 and has a basis of elements with infinite order; Theorem 1.9 now yields that F is paradoxical. Let

$$D = \{s \in S^2 \mid \exists \sigma \in F \backslash \{1\} : \sigma(s) = s\}.$$

Since $F\backslash\{1\}$ only contains rotations about axes through the origin, each $\sigma \in F\backslash\{1\}$ has exactly two fixed points in S^2. As F is countable, we conclude that D is countable too. Now, since the action of F on $S^2\backslash D$ is without nontrivial fixed points, Theorem 1.4 shows that $S^2\backslash D$ is F-paradoxical and therefore SO_3-paradoxical.

□

Recalling the free product from the previous chapter, we see that F actually is the free product $\mathbb{Z} * \mathbb{Z}$ and that T is $\mathbb{Z}_2 * \mathbb{Z}_3$, since the free groups have pairwise "isomorphic" bases (in the sense discussed on page 4). Therefore, the above asserts that the paradoxical group $\mathbb{Z} * \mathbb{Z}$ is a subgroup of $\mathbb{Z}_2 * \mathbb{Z}_3$, and Corollary 1.5 says that $\mathbb{Z}_2 * \mathbb{Z}_3$ is paradoxical as well.

2.2 Equidecomposability

The Hausdorff paradox ensures that S^2 is almost paradoxical with respect to SO_3. To improve this result, we introduce the concept of equidecomposability, which divides $\mathcal{P}(\mathbb{R}^3)$ into equivalence classes. We then prove that paradoxicality is a class property and, finally, show that $S^2\backslash D$ from Theorem 2.2 and S^2 belong to the same class.

Definition 2.3 *Suppose the group G acts on a set X. Two subsets $A, B \subseteq X$ are said to be G-*equidecomposable*, and we write $A \sim_G B$, if there exists a partitioning $A_1, \ldots, A_n$ of A and a partitioning $B_1, \ldots, B_n$ of B and elements $g_1, \ldots, g_n \in G$, such that $g_i.A_i = B_i$ for $i = 1, \ldots, n$.*

Note that the function $g \colon A \to B$, given by

$$g(x) = \begin{cases} g_1.x & \text{for } x \in A_1 \\ \vdots & \vdots \\ g_n.x & \text{for } x \in A_n \end{cases} \tag{2.1}$$

is a bijection. On the other hand, if there exists a bijection in the form (2.1), then $A \sim_G B$. Hence, equidecomposability can also be defined as piecewise G-congruence. We will repeatedly make use of this alternative definition.

Proposition 2.4 *Equidecomposability is an equivalence relation on $\mathcal{P}(X)$.*

PROOF: It is obvious that $\sim_G$ is reflexive and symmetric. To show that $\sim_G$ is transitive, let $A_1, \ldots, A_n \subseteq A$, $B_1, \ldots, B_n \subseteq B$ and $g_1, \ldots, g_n \in G$ demonstrate that $A \sim_G B$, and let $B'_1, \ldots, B'_m \subseteq B$, $C_1, \ldots, C_m \subseteq C$ and $h_1, \ldots, h_m \in G$ demonstrate that $B \sim_G C$. For $i = 1, \ldots, n$ and $j = 1, \ldots, m$, define

$$B_{ij} = B_i \cap B'_j.$$

First assume that $B_{ij} \neq \varnothing$ for all i, j. Then, since B_{ij} partition B, $A_{ij} = g_i^{-1}.B_{ij}$ and $C_{ij} = h_j.B_{ij}$ are partitionings of A and C respectively, and

$$C_{ij} = (h_j g_i).A_{ij}.$$

Hence, $A \sim_G C$. If some of the B_{ij}'s are empty, we can ignore them, obtain a partitioning in less than nm sets and then use the same argument. In all cases, $A \sim_G C$, and it follows that $\sim_G$ is an equivalence relation on $\mathcal{P}(X)$.

□

The transitivity of $\sim_G$ shows that whenever the bijections g and h are in the form (2.1), so is the bijection $h \circ g$. This will be used repeatedly hereafter without further discussion.

The following theorem reformulates Definition 1.1 in terms of equidecomposability.

Theorem 2.5 *A subset $E \subseteq X$ is G-paradoxical if, and only if, there exist disjoint sets $A, B \subseteq E$, such that $A \sim_G E$ and $B \sim_G E$.*

PROOF: It is trivial to see that, if $A \sim_G E$ and $B \sim_G E$ for some disjoint subsets $A, B \subseteq E \subseteq X$, then E is G-paradoxical. Conversely, assume that $E \subseteq X$ is G-paradoxical, that is, there exist pairwise disjoint sets $A_1, \ldots, A_n$,$B_1, \ldots, B_m$ of E and elements $g_1, \ldots, g_n$,$h_1, \ldots, h_m \in G$ with $E = \bigcup_{i_1}^{n} g_i.A_i = \bigcup_{j_1}^{m} h_j.B_j$. As the unions are not necessarily over disjoint sets, we define $A'_1 = A_1$ and

$$A'_k = A_k \backslash g_k^{-1}.(\bigcup_{i=1}^{k-1} g_i.A'_i)$$

for $k = 2, \ldots, n$. We now find that $E = \bigcup_{i=1}^{n} g_i.A'_i$ is a disjoint union. Hence, $E \sim_G A' = \bigcup_{i=1}^{n} A'_i$. Similarly, we can define B' and obtain $E \sim_G B'$. Since $A' \subseteq A$ and $B' \subseteq B$, where $A = \bigcup_{i=1}^{n} A_i$ and $B = \bigcup_{j=1}^{m} B_j$, we find that $A' \cap B' = \varnothing$ as required.

□

Now, with equidecomposability linked to paradoxicality, we show the following essential result.

Proposition 2.6 *Paradoxicality is a class property with respect to equidecomposability.*

Proof: Suppose G acts on X and E, E' are G-equidecomposable subsets of X, and that E is G-paradoxical. We want to show that E' is G-paradoxical. According to Theorem 2.5, let $A, B \subseteq E$ be two disjoint subsets with $A \sim_G E \sim_G B$. Furthermore, let $E_1, \ldots, E_n \subseteq E$, $E'_1, \ldots, E'_n \subseteq E'$ and $g_1, \ldots, g_n \in G$ witness that $E \sim_G E'$. Now, define

$$A' = \bigcup_{i=1}^{n} g_i.(E_i \cap A) \quad \text{and} \quad B' = \bigcup_{i=1}^{n} g_i.(E_i \cap B).$$

Since $\bigcup_{i=1}^{n}(E_i \cap A) = A$, it follows that $A \sim_G A'$ and equivalently that $B \sim_G B'$. As A and B are disjoint, then so are A' and B'. We now have two disjoint subsets $A', B' \subseteq E'$ with $A' \sim_G A \sim_G E \sim_G E'$ and $B' \sim_G B \sim_G E \sim_G E'$, and due to transitivity, $A' \sim_G E' \sim_G B'$. Hence E' is G-paradoxical.

□

2.3 The Banach–Tarski Paradox

Now we are ready to expand the Hausdorff paradox to the entire unit sphere S^2. The idea of the proof is to *absorb* the points in D; proofs by absorption are standard tricks in the theory of equidecomposability.

Lemma 2.7 *If D is a countable subset of S^2, then $S^2 \sim_{\mathrm{SO}_3} S^2 \backslash D$.*

Proof: Let l be a line through the origin, such that $l \cap D = \varnothing$. (The line l can be found, since D is only countable.) Define

$$A = \{\theta \in [0, 2\pi[\mid \exists n \in \mathbb{N}\ \exists P \in D : \rho_\theta^n(P) \in D\},$$

where ρ_θ is the rotation with angle θ about l. For fixed $P \in D$, there are only countably many θ such that $\rho_\theta(P) \in D$. Since $\rho_\theta^n(P) \in D$ if, and only if, $\rho_{n\theta}(P) \in D$, there are only countably many (n, θ) such that $\rho_\theta^n(P) \in D$. Since D is countable it follows that A is countable. We can therefore choose $\theta_0 \in [0, 2\pi[\backslash A$ and find that $D, \rho_{\theta_0}(D), \rho_{\theta_0}^2(D), \ldots$ are pairwise disjoint (since $\rho_{\theta_0}^n(D) \cap \rho_{\theta_0}^m(D) \neq \varnothing$ implies $\rho_{\theta_0}^{n-m}(D) \cap D \neq \varnothing$, which contradicts the choice of θ_0). Letting $\widehat{D} = \bigcup_{n=0}^{\infty} \rho_{\theta_0}^n(D)$, we see that

$$S^2 = \widehat{D} \cup (S^2 \backslash \widehat{D}) \sim_{\mathrm{SO}_3} \rho_{\theta_0}(\widehat{D}) \cup (S^2 \backslash \widehat{D}) = S^2 \backslash D,$$

as required.

□

Together with the Hausdorff paradox and Proposition 2.6, Lemma 2.7 shows that S^2 is SO_3-paradoxical. An easy consequence of this is the Banach–Tarski paradox.

Theorem 2.8 (The Banach–Tarski Paradox)(AC) *The unit ball B in $\mathbb{R}^3$ is SO_3-paradoxical.*

Proof: The preceding comments and the radial correspondence between S^2 and $B\backslash\{0\}$, given by $P \mapsto \{rP \mid 0 < r \leq 1\}$, shows that $B\backslash\{0\}$ is SO_3-paradoxical. To complete the proof, we only need to *absorb* the point 0, that is, show that $B \sim_{\mathrm{SO}_3} B\backslash\{0\}$. Let l denote the line through $(0,0,\frac{1}{2})$ and $(0,1,\frac{1}{2})$ and choose a rotation ρ of infinite order about l. Now, as in the proof of Lemma 2.7, let $\widehat{D} = \{\rho^n(0) \mid n \in \mathbb{N}_0\}$. Since $\rho(\widehat{D}) = \widehat{D}\backslash\{0\}$, $B \sim_{\mathrm{SO}_3} B\backslash\{0\}$, as required.

□

By letting $0 < r \leq R$ in the radial correspondence, the Banach–Tarski paradox can easily be expanded to cover all solid balls centred at the origin and with radius R. Since all solid balls of radius R in $\mathbb{R}^3$ are O_3-equidecomposable[1] and $\mathrm{SO}_3 \subseteq \mathrm{O}_3$, it follows that any solid ball is O_3-paradoxical. By letting $0 < r < \infty$, we can even find that $\mathbb{R}^3\backslash\{0\}$ is SO_3-paradoxical, and since the origin can be absorbed, $\mathbb{R}^3$ is SO_3-paradoxical. Actually, we can improve this even more by the remarking result that any two bounded subsets with nonempty interiors are O_3-equidecomposable. To prove this we introduce an ordering $\preccurlyeq_{\mathrm{O}_3}$ on $\mathcal{P}(\mathbb{R}^3)$.

Definition 2.9 *For $A, B \subseteq \mathbb{R}^3$, $A \preccurlyeq_G B$ if A is G-equidecomposable with a subset of B.*

The relation $\preccurlyeq_G$ is actually a partial ordering of the $\sim_G$-equivalence classes in $\mathcal{P}(X)$. The difficult part in proving this is showing that $\preccurlyeq_G$ is antisymmetrical. This is contained in the following generalization of Bernstein's equivalence theorem.

Theorem 2.10 (Banach–Schröder–Bernstein) *Suppose a group G acts on the set X, and A, B are subsets of X. If $A \preccurlyeq_G B$ and $B \preccurlyeq_G A$, then $A \sim_G B$.*

Proof: The assumptions ensure the existence of subsets $A' \subseteq A$ and $B' \subseteq B$ with $A \sim_G B'$ and $A' \sim_G B$. Let $f\colon A \to B'$ and $g\colon A' \to B$ be bijections as in (2.1). Note that since the restriction of f to $C \subseteq A$ is written in the form (2.1) as well, we have that $C \sim_G f(C)$, whenever $C \subseteq A$. Now, let $C_0 = A\backslash A'$ and inductively define

$$C_{i+1} = (g^{-1} \circ f)(C_i) \quad \text{for } i = 0, 1, \ldots.$$

[1] O_n denotes the set of isometries in $\mathbb{R}^n$.

Letting $C = \bigcup_{i=0}^{\infty} C_i$, we obtain $g(A\backslash C) = B\backslash f(C)$, since $A\backslash C = A'\backslash C = A'\backslash((g^{-1} \circ f)(C) \cup C_0) = A'\backslash(g^{-1} \circ f)(C)$. Hence, $A\backslash C \sim_G B\backslash f(C)$, and since $C \sim_G f(C)$ as mentioned, we conclude that $(A\backslash C) \cup C \sim_G (B\backslash f(C)) \cup f(C)$, that is, $A \sim_G B$.

□

Since the reflexivity and transitivity of $\preccurlyeq_G$ is obvious and $\preccurlyeq_G$ clearly is a class property with respect to equidecomposability, an immediate consequence of Theorem 2.10 is:

Corollary 2.11 *The relation $\preccurlyeq_G$ is an ordering of the $\sim_G$-classes in $\mathcal{P}(X)$.*

We can now generalize the Banach–Tarski paradox to cover all bounded subsets of $\mathbb{R}^3$ with a nonempty interior.

Theorem 2.12 (The Banach–Tarski Paradox – strong form)(AC) *If A and B are two bounded subsets of $\mathbb{R}^3$, each having a nonempty interior, then $A \sim_{\mathrm{O}_3} B$.*

Proof: It suffices to show that $A \preccurlyeq_{\mathrm{O}_3} B$, since a similar argument yields $B \preccurlyeq_{\mathrm{O}_3} A$, and the theorem then follows from the antisymmetry of $\preccurlyeq_{\mathrm{O}_3}$. Choose solid balls, $K, L \subseteq \mathbb{R}^3$, such that $A \subseteq K$ and $L \subseteq B$. Let $S = L_1 \cup \cdots \cup L_N$ be the set consisting of N disjoint copies of L, where N is chosen large enough that we can find $t_1, \ldots, t_N \in \mathrm{O}_3$ such that $K \subseteq \bigcup_{i=1}^{N} t_i.L_i$. (This is possible by the boundedness of K.) By repeatedly applying the original version of the Banach–Tarski paradox, the remarks following this show that $S \sim_{\mathrm{O}_3} L$. As there obviously exists an injection $K \to S$ in the form (2.1), $A \subseteq K \preccurlyeq_{\mathrm{O}_3} S \sim_{\mathrm{O}_3} L \subseteq B$, so $A \preccurlyeq_{\mathrm{O}_3} B$ as required.

□

Chapter 3

The Type Semigroup and Tarski's Theorem

3.1 The Type Semigroup

If a group G acts on a set X, we can let $X^* = X \times \mathbb{N}_0$ and $G^* = G \times \mathrm{Perm}(\mathbb{N}_0)$, and define an enlarged action of G^* on X^* by $(g, \sigma).(x, n) = (g.x, \sigma(n))$. The new set X^* equips us with an infinite multitude of copies of X and enables us to define G-paradoxicality of a subset $E \subseteq X$ as $E \times \{0\} \sim_{G^*} E \times \{0, 1\}$, since $E_1 \sim_G E_2$ if, and only if, $E_1 \times \{n\} \sim_{G^*} E_2 \times \{m\}$ for all $E_1, E_2 \subseteq X$ and $n, m \in \mathbb{N}_0$.

Definition 3.1 *If $A \subseteq X^*$, then the elements in $\pi_2(A)$ (where π_2 is the projection of X^* on $\mathbb{N}_0$) are called the* levels *of A. A is called* bounded *if it has only finitely many levels.*

Proposition 2.4 shows that $\sim_{G^*}$ defines an equivalence relation on $\mathcal{P}(X^*)$. If a subset $A \subseteq X^*$ is bounded and $A \sim_{G^*} B$, then B is bounded as well. We can therefore define as follows.

Definition 3.2 *The equivalence class with respect to G^*-equidecomposability of a bounded subset $A \subseteq X^*$ is called the* type *of A and is denoted $[A]$. The collection of types of bounded sets is denoted $\mathcal{S}$.*

For a bounded set $B \subseteq X^*$, an upward shift of B is a set

$$B' = \{(b, n + k) \in X^* \mid (b, n) \in B\},$$

where $k \in \mathbb{N}$. We can now define the composition $+$ on $\mathcal{S}$ by $[A] + [B] = [A \cup B']$, where B' is an upward shift of B, such that the levels of B' are disjoint from the levels of A. Obviously, the composition is independent of the choice of representatives and is therefore welldefined. Furthermore $+$ is commutative and associative, so $(\mathcal{S}, +)$ is an abelian semigroup with identity $[\varnothing]$. We shall refer to $(\mathcal{S}, +)$ as the type semigroup with respect to G and X. The easiest way to show that two types are identical is often to make use of the antisymmetry of the ordering $\leq$, given by $\alpha \leq \beta$ for $\alpha, \beta \in \mathcal{S}$ if, and only if, there exists $\gamma \in \mathcal{S}$ such that $\alpha + \gamma = \beta$. To realize that $\leq$ actually defines an ordering on $\mathcal{S}$, one can run through the proof of the Banach–Schröder–Bernstein theorem.

The ordered type semigroup $(\mathcal{S}, +, \leq)$ satisfies many simple properties. Mostly we are interested in the apparently trivial cancellation law: $n\alpha = n\beta \Rightarrow \alpha = \beta$. This, however, is based on König's theorem of graph theory. Therefore, recall that a graph is a set V of vertices with a collection E of edges consisting of unordered pairs of distinct elements of V. We allow graphs to have infinitely many vertices and multiple edges. A graph is called bipartite if the vertex set splits into two pieces so that each edge has one vertex in each piece. The degree of a vertex v is the number of edges containing v. A graph is called k-regular if all vertices in the graph have degree k. Finally, a perfect matching M is a subset of E, satisfying the condition that each vertex $v \in V$ is contained in one, and only one, edge in M.

Theorem 3.3 (König's Theorem)(AC) *A k-regular bipartite graph (V, E), where $k < \infty$ has a perfect matching.*

PROOF: Let $u \leftrightarrow v$ if there exists a finite path from u to v. Since each edge is unordered, $\leftrightarrow$ divides V into connected equivalence classes, called components; in the infinite case, this requires the axiom of choice, since we need to select a representative from each class. Note that a perfect matching on each of the components can be united to a perfect matching on the entire graph. Thus, it suffices to find a perfect matching on each component. As the graph is k-regular, the number of vertices that can be reached from a given vertex in a path of length n is bounded by k^n. Hence, each component is at most countable, so we need only consider connected graphs that are finite or countably infinite.

So assume that $\mathcal{G} = (V, E)$ is a *finite*, connected, bipartite and k-regular graph. The bipartiteness of $\mathcal{G}$ ensures that $V = \{a_1, \ldots, a_n, b_1, \ldots, b_n\}$ and that E consists only of edges written as $a_i b_j$. (Note that, since $\mathcal{G}$ is k-regular, there must be as many a_i's as there are b_j's.) By induction on n, we want to show that G has a perfect matching, the case $n = 1$ being trivial. So, assume that $\mathcal{G}$ has a perfect matching for all $n < m$ and consider the case $n = m$. Transform $\mathcal{G}$ by the following algorithm:

(i) Choose j so that $a_1 b_j$ is in E and remove (one of) the edge(s) $a_1 b_j$. Let $x = b_j$.

(ii) If $x = b_j$, which it has just entered from a_s, choose the lowest $i \neq s$ such that $a_i b_j$ is in E. Add an additional edge $a_i b_j$ to E and set $x = a_i$.

(iii) If $x = a_i$, which it has just entered from b_t, choose the lowest $j \neq t$ such that $a_i b_j$ is in E. Remove (one of) the edge(s) $a_i b_j$ and set $x = b_j$. Go to (ii).

The algorithm terminates as soon as it is impossible to find an i in (ii) or a j in (iii). Note that all edges arising in the transformed graphs are parallel to edges in $\mathcal{G}$. We suggest that the algorithm will terminate at some point, ending with $x = a_1$.

To see this, first suppose that the algorithm stops at $x = b_j$ for some j. If $x = a_i$ in the previous step, the edge $a_i b_j$ has to have a multiplicity of k at the point before termination. Since the algorithm cannot stop in step two (as this would contradict the connectedness of $\mathcal{G}$), we can always find $b_t \neq b_j$ such that x came to a_i from b_t. Therefore, at the point before termination, (one of) the edge(s) $a_i b_t$ was doubled, which leaves a_i with a degree of at least $k+2$, which is impossible.

Next, suppose that the algorithm terminates with $x = a_i$ where $i \neq 1$. If $x = b_j$ in the previous step, the edge $a_i b_j$ has to have multiplicity $k+1$ at point of termination. Consequently, b_j has degree at least $k+1$ at the point of termination, and this is impossible.

Finally, suppose that the algorithm never terminates. First, note that since the degrees of all vertices are bounded by $k-1$ and $k+1$, an edge passed infinitely often must be passed infinitely often in both directions. Consider a vertex b_j and suppose that three edges $a_i b_j$, $a_l b_j$ and $a_s b_j$, where $i < l < s$, are passed infinitely often. Then the edge $a_s b_j$ can only be passed in the direction from a_s to b_j, since $a_i b_j$ or $a_l b_j$ is always preferred by the algorithm, and this is impossible. Similarly, each of the a_i's are endpoints of at most two infinitely used edges. At the point in the algorithm where all finitely used edges have already been passed, x therefore moves around in a cycle. This, however, contradicts the proposition that all the edges are passed infinitely often in both directions.

Based on this, we conclude that the algorithm at some point terminates with $x = a_1$. The construction of the algorithm now implies that a_1 in the terminating graph is connected to some b_j by k parallel edges. This reveals a graph with k-regular and bipartite components, each having less than $2m$ elements. Using the induction hypothesis on these components, we obtain the desired perfect matching in the case $n = m$, and the result follows by induction.

Next, assume that $\mathcal{G}$ is a *countably infinite*, connected, bipartite and k-regular graph. Note that, since $\mathcal{G}$ is k-regular, both V and E are countably infinite, so we can enumerate E by $E = \{e_n \mid n \in \mathbb{N}\}$. Consider the collection of all finite sequences in $\{0, 1\}$. Such a sequence s of length n is said to be good if there exists a finite k-regular graph (V_n, E_n) with $\{e_1, \ldots, e_n\} \subseteq E_n$ and a perfect matching $M_n \subseteq E_n$, such that $e_i \in M_n$ if, and only if, 1 is the ith element of s. (Observe that the condition $\{e_1, \ldots, e_n\} \subseteq E_n$ implies that the vertices appearing in

$\{e_1, \ldots, e_n\}$ are contained in V_n.)

First, let us realize that, for all $n \in \mathbb{N}$, there is at least one good sequence of length n. Consider the finite bipartite graph, consisting of the edges $e_1, \ldots, e_n$ and the vertices emanating from these. Adding vertices to equalize the two parts of the graph and adding edges wherever necessary to push the degree of each vertex up to k produces a finite, bipartite and k-regular graph (V_n, E_n). By the finite case, (V_n, E_n) has a perfect matching M_n, so the sequence $s = (1_{M_n}(e_1), \ldots, 1_{M_n}(e_n))$ is good as required.

Based on this, we conclude that there are infinitely many good sequences. Since there are only finitely many sequences in $\{0, 1\}$ of length i, we can therefore inductively define the sequence $s = (s_i)_{i \in \mathbb{N}}$, satisfying the condition that $(s_1, \ldots, s_i)$ has infinitely many good extensions for $i = 1, 2, \ldots$. Note that each initial segment of s is good, since any segment with a good extension is necessarily good itself. We suggest that $M = \{e_i \mid s_i = 1\}$ is a perfect matching of $\mathcal{G}$. To see that any vertex $v \in V$ appears in M, choose n sufficiently large that all k edges emanating from v are contained in $\{e_1, \ldots, e_n\}$. As the initial segment $(s_1, \ldots, s_n)$ is good, there exists a finite, bipartite and k-regular graph (V_n, E_n) with $\{e_1, \ldots, e_n\} \subseteq E_n$ and a perfect matching M_n. Hence, there exists an $i \leq n$ such that $s_i = 1$ and e_i is an edge emanating from v, so v appears in M. Clearly v cannot appear twice in M, since this would contradict M_n being a perfect matching on (V_n, E_n). Consequently, M is the required perfect matching on $\mathcal{G}$, and the result follows.

□

After this rather lengthy proof, let us return to the type semigroup $(\mathcal{S}, +)$ with respect to G and X and state the essential cancellation law.

Theorem 3.4 (AC) *For $\alpha, \beta \in \mathcal{S}$ and $n \in \mathbb{N}$, $n\alpha = n\beta$ implies $\alpha = \beta$.*

Proof: Suppose $n\alpha = n\beta$, that is, there exist two disjoint, bounded and G^*-equidecomposable subsets $E, E' \subseteq X^*$, with partitionings $E = \bigcup_{i=1}^n A_i$ and $E' = \bigcup_{j=1}^n B_j$ such that $[A_i] = \alpha$ and $[B_j] = \beta$ for all i, j. Let $\chi \colon E \to E'$ be in the form (2.1), demonstrating that $E \sim_{G^*} E'$, and let likewise $\phi_i \colon A_1 \to A_i$ and $\psi_j \colon B_1 \to B_j$ demonstrate that $A_1 \sim_{G^*} A_i$ and $B_1 \sim_{G^*} B_j$, taking ϕ_1 and ψ_1 to be the identity. For all $a \in A_1$ and $b \in B_1$, define

$$\bar{a} = \{a, \phi_2(a), \ldots, \phi_n(a)\} \quad \text{and} \quad \bar{b} = \{b, \psi_2(b), \ldots, \psi_n(b)\}.$$

Note that, since ϕ_i and ψ_j are bijections, $\{\bar{a} \mid a \in A_1\}$ and $\{\bar{b} \mid b \in B_1\}$ are partitionings of E and E' respectively.

Now, we can form a bipartite graph by letting $\{\bar{a} \mid a \in A_1\}$ and $\{\bar{b} \mid b \in B_1\}$ be the two parts of the vertex set, and, for each $i = 1, \ldots n$, introducing the edge $\bar{a}\bar{b}$ if $\chi(\phi_i(a)) \in \bar{b}$. Obviously, each $\bar{a}$ has degree n (counted with multiplicity). However, each $\bar{b}$ appears in the edges corresponding to $\chi^{-1}(\psi_j(b))$ for $j = 1, \ldots n$.

Hence, the graph is n-regular, and König's theorem ensures that it has a perfect matching M. For each $\bar{a}$, we can now uniquely determine $\bar{b}$, such that $\bar{a}\bar{b} \in M$, so

$$C_{ij} = \{a \in A_1 \mid \bar{a}\bar{b} \in M \Rightarrow \chi(\phi_i(a)) = \psi_j(b)\}$$

is a partitioning of A_1. Similarly,

$$D_{ij} = \{b \in B_1 \mid \bar{a}\bar{b} \in M \Rightarrow \chi(\phi_i(a)) = \psi_j(b)\}$$

is a partitioning of B_1. Furthermore, the mapping $\psi_j^{-1} \circ \chi \circ \phi_i \colon C_{ij} \to D_{ij}$ is a bijection in the form (2.1), since ψ_j, χ and ϕ_i are, so the bijection $A_1 \to B_1$ given by $x \mapsto (\psi_j^{-1} \circ \chi \circ \phi_i)(x)$ for $x \in C_{ij}$ is also in the form (2.1). Hence, $A_1 \sim_{G^*} B_1$, or $\alpha = \beta$, as desired.

□

Since $E \subseteq X$ is G-paradoxical if, and only if, $[E] = 2[E]$, we intend to make use of the following corollary.

Corollary 3.5 (AC) *If $\alpha \in \mathcal{S}$ and $n \in \mathbb{N}_0$ satisfy $(n+1)\alpha \leq n\alpha$, then $\alpha = 2\alpha$.*

Proof: Successive use of the hypothesized inequality yields for $n \geq 1$ that $n\alpha \geq (n+1)\alpha = n\alpha + \alpha \geq \cdots \geq n\alpha + n\alpha = 2n\alpha$. Obviously $n\alpha \leq 2n\alpha$, so by the antisymmetry of $\leq$, $n\alpha = 2n\alpha$, and the cancellation law gives $\alpha = 2\alpha$. The case $n = 0$ is trivial.

□

3.2 Tarski's Theorem

Tarski's theorem ties paradoxicality and measure theory together by stating that a subset E of X is not G-paradoxical if, and only if, there exists a finitely additive and G-invariant measure μ on $(X, \mathcal{P}(X))$, normalizing E. The latter implies the former, since the paradoxicality of E gives $2 = 2\mu(E) \leq \mu(E) = 1$. To prove the other direction, we shall attempt to construct a measure ν on $\mathcal{S}$ (that is, a mapping $\nu \colon \mathcal{S} \to [0, \infty]$ with $\nu(\alpha+\beta) = \nu(\alpha)+\nu(\beta)$), such that $\nu([E \times \{0\}]) = 1$. The measure ν induces a measure μ on $(X, \mathcal{P}(X))$ by letting $\mu(A) = \nu([A \times \{0\}])$, and μ obviously has the desired properties.

In the following lemma and theorem, we consider a commutative semigroup $(\mathcal{T}, +)$ on which we can define the preordering $\alpha \leq \beta$ if $\alpha + \gamma = \beta$ for some $\gamma \in \mathcal{T}$. Note that the ordering is not necessarily total. For $\varepsilon \in \mathcal{T}$, we say that α is bounded with respect to ε, if there exists $n \in \mathbb{N}$, such that $\alpha \leq n\varepsilon$.

Lemma 3.6 *Suppose $(\mathcal{T}, +)$ is a commutative semigroup and that $\varepsilon \in \mathcal{T}$. Furthermore, assume that all elements of $\mathcal{T}$ are bounded by ε and that $(n+1)\varepsilon \nleq n\varepsilon$*

for all $n \in \mathbb{N}_0$*. If* $\mathcal{T}_0$ *is a finite subset of* $\mathcal{T}$ *and* $\varepsilon \in \mathcal{T}_0$*, then there exists a function* $\nu\colon \mathcal{T}_0 \to [0,\infty]$*, such that* $\nu(\varepsilon) = 1$ *and* $\nu(\alpha+\beta) = \nu(\alpha)+\nu(\beta)$ *whenever* $\alpha, \beta, \alpha+\beta \in \mathcal{T}_0$*.*

PROOF: Instead of proving the existence of a measure as above, we will, under the hypothesized conditions, prove the existence of a measure $\nu\colon \mathcal{T}_0 \to [0,\infty]$ with $\nu(\varepsilon) = 1$ and $\sum_{i=1}^n \nu(\alpha_i) \leq \sum_{j=1}^m \nu(\beta_j)$ for all α_i, β_j, satisfying $\alpha_1 + \cdots + \alpha_n \leq \beta_1 + \cdots + \beta_m$. This is sufficient, since the additive property of ν easily follows from the last of these conditions.

The proof is by induction on the size of $\mathcal{T}_0$. The case $|\mathcal{T}_0| = 1$ implies $\mathcal{T}_0 = \{\varepsilon\}$, so $\nu(\varepsilon) = 1$ is the desired function. The fact that ν satisfies the latter condition can be reduced to showing that $m\varepsilon \leq n\varepsilon$ imply $m \leq n$, but this follows from the hypothesis on ε, since $m\varepsilon \leq n\varepsilon$ and $m \geq n+1$ implies $(n+1)\varepsilon \leq m\varepsilon \leq n\varepsilon$, which is a contradiction.

Suppose next that $|\mathcal{T}_0| > 1$ and that ν exists in any smaller case of $|\mathcal{T}_0|$. Let ξ be an element in $\mathcal{T}_0 \backslash \{\varepsilon\}$ and let ν' be a function on $\mathcal{T}_0 \backslash \{\xi\}$ satisfying the conditions of the lemma. Define $\nu\colon \mathcal{T}_0 \to [0,\infty]$ by

$$\nu(\tau) = \begin{cases} \nu'(\tau) & \text{for } \tau \in \mathcal{T}_0 \backslash \{\xi\} \\ \inf\{\frac{\sum_{i=1}^p \nu'(\delta_i) - \sum_{j=1}^q \nu'(\gamma_j)}{r}\} & \text{for } \tau = \xi \end{cases} \tag{3.1}$$

where the infimum is taken over all $r \in \mathbb{N}$ and $\delta_i, \gamma_j \in \mathcal{T}_0 \backslash \{\xi\}$, satisfying the condition that $\gamma_1 + \cdots + \gamma_q + r\xi \leq \delta_1 + \cdots + \delta_p$. Since all elements in $\mathcal{T}_0 \backslash \{\xi\}$ are bounded by some $n\varepsilon$, ν' takes only finite values, so the set on which the infimum is taken is welldefined. That ν actually maps $\mathcal{T}_0$ to $[0,\infty]$ is a consequence of the desired property of ν, since $\varepsilon \leq \varepsilon + \xi$ then implies $1 \leq 1 + \nu(\xi)$. It therefore only remains to show that ν satisfies the latter condition.

We want to show that, if $\alpha_1 + \cdots + \alpha_m + s\xi \leq \beta_1 + \cdots + \beta_n + t\xi$, where $\alpha_i, \beta_j \in \mathcal{T}_0 \backslash \{\xi\}$ and $s, t \in \mathbb{N}_0$, then $\sum_{i=1}^m \nu'(\alpha_i) + s\nu(\xi) \leq \sum_{j=1}^n \nu'(\beta_j) + t\nu(\xi)$. We divide into three possible cases.

In the case $s = 0$ and $t = 0$, the desired inequality follows from the properties of ν'.

Next, assume $s = 0$ and $t > 0$. We want to obtain that

$$\nu(\xi) \geq \frac{\sum_{i=1}^m \nu'(\alpha_i) - \sum_{j=1}^n \nu'(\beta_j)}{t}.$$

It suffices to show that the following inequality holds for all δ_i, γ_j, r satisfying the conditions of the infimum in (3.1):

$$\frac{\sum_{i=1}^p \nu'(\delta_i) - \sum_{j=1}^q \nu'(\gamma_j)}{r} \geq \frac{\sum_{i=1}^m \nu'(\alpha_i) - \sum_{j=1}^n \nu'(\beta_j)}{t}. \tag{3.2}$$

From the given inequality $\alpha_1 + \cdots + \alpha_m \leq \beta_1 + \cdots + \beta_n + t\xi$, we obtain

$$\begin{aligned} r\alpha_1 + \cdots + r\alpha_m + t\gamma_1 + \cdots + t\gamma_q &\leq r\beta_1 + \cdots + r\beta_n + rt\xi + t\gamma_1 + \cdots + t\gamma_q \\ &\leq r\beta_1 + \cdots + r\beta_n + t\delta_1 + \cdots + t\delta_p. \end{aligned}$$

The induction hypothesis on ν' now yields

$$r\sum_{i=1}^{m}\nu'(\alpha_i)+t\sum_{j=1}^{q}\nu'(\gamma_j)\le r\sum_{j=1}^{n}\nu'(\beta_j)+t\sum_{i=1}^{p}\nu'(\delta_i),$$

which implies the desired inequality (3.2).

Finally, suppose $s>0$. It suffices to show that

$$\sum_{i=1}^{m}\nu'(\alpha_i)+s\nu(\xi)\le\sum_{j=1}^{n}\nu'(\beta_j)+tz,$$

where

$$z=\frac{\sum_{i=1}^{p}\nu'(\delta_i)-\sum_{j=1}^{q}\nu'(\gamma_j)}{r}$$

for arbitrary γ_j,δ_i,r satisfying the conditions of the infimum in (3.1). From the given inequality $\alpha_1+\cdots+\alpha_m+s\xi\le\beta_1+\cdots+\beta_n+t\xi$, we obtain

$$\begin{aligned} r\alpha_1+\cdots+r\alpha_m+t\gamma_1+\cdots+t\gamma_q+rs\xi &\le r\beta_1+\cdots+r\beta_n+rt\xi+t\gamma_1+\cdots+t\gamma_q\\ &\le r\beta_1+\cdots+r\beta_n+t\delta_1+\cdots+t\delta_p. \end{aligned}$$

If the second step is ignored, the above inequality is a typical one used to define $\nu(\xi)$, and it follows that

$$\begin{aligned} \sum_{i=1}^{m}\nu'(\alpha_i)+s\nu(\xi) &\le \sum_{i=1}^{m}\nu'(\alpha_i)+\\ &\quad s\frac{r\sum_{j=1}^{n}\nu'(\beta_j)+t\sum_{i=1}^{p}\nu'(\delta_i)-r\sum_{i=1}^{m}\nu'(\alpha_i)-t\sum_{j=1}^{q}\nu'(\gamma_j)}{rs}\\ &= \sum_{j=1}^{n}\nu'(\beta_j)+tz, \end{aligned}$$

as required. Hence, ν has the desired properties.

□

Theorem 3.7 (AC) *Suppose $(\mathcal{T},+)$ is a commutative semigroup and that $\varepsilon\in\mathcal{T}$. Then the following statements are equivalent:*

(i) *For all $n\in\mathbb{N}_0$, $(n+1)\varepsilon\not\le n\varepsilon$.*

(ii) *There exists a finitely additive measure μ on $\mathcal{T}$ (that is, a function $\mu\colon\mathcal{T}\to[0,\infty]$ with $\mu(\alpha+\beta)=\mu(\alpha)+\mu(\beta)$ for all $\alpha,\beta\in\mathcal{T}$), such that $\mu(\varepsilon)=1$.*

Proof: Suppose $(n+1)\varepsilon \leq n\varepsilon$, that is, there exists $\gamma \in \mathcal{T}$, such that $(n+1)\varepsilon + \gamma = n\varepsilon$. Then there cannot exist a function μ satisfying the conditions in (ii), as this would require $n+1 = \mu((n+1)\varepsilon) \leq \mu((n+1)\varepsilon + \mu(\gamma) = \mu(n\varepsilon) = n$. Hence, (ii) implies (i).

To prove that (i) implies (ii), we can, without loss of generality, assume that all elements of $\mathcal{T}$ are bounded with respect to ε, for once we have a measure on the bounded elements, it can be extended by assigning the unbounded elements measure ∞. So, suppose that $(n+1)\varepsilon \nleq n\varepsilon$ for all $n \in \mathbb{N}$, and, for any finite subset $\mathcal{T}_0 \subseteq \mathcal{T}$ containing ε, let $\mathcal{M}(\mathcal{T}_0)$ consist of all functions $f \in [0,\infty]^{\mathcal{T}}$, satisfying the condition that $f(\varepsilon) = 1$ and $f(\alpha+\beta) = f(\alpha) + f(\beta)$ whenever $\alpha, \beta, \alpha+\beta \in \mathcal{T}_0$. Note, that according to Lemma 3.6, $\mathcal{M}(\mathcal{T}_0)$ is nonempty.

Since $[0,\infty]^{\mathcal{T}}$ is a product of compact spaces, Tychonoff's theorem (which requires the axiom of choice) shows that $[0,\infty]^{\mathcal{T}}$ is compact. We therefore know that the intersection of a collection of closed subsets of $[0,\infty]^{\mathcal{T}}$ is nonempty whenever any intersection of finitely many members of the collection is nonempty. Therefore, if we can prove that the $\mathcal{M}(\mathcal{T}_0)$'s are closed and that any intersection of finitely many of them is nonempty, we can conclude that there exists a measure μ that lies in every $\mathcal{M}(\mathcal{T}_0)$. This measure is as desired, since $\mu(\alpha+\beta) = \mu(\alpha)+\mu(\beta)$ follows from the fact that $\mu \in \mathcal{M}(\{\{\varepsilon, \alpha, \beta, \alpha+\beta\})$.

To prove that $\mathcal{M}(\mathcal{T}_0)$ is closed in $[0,\infty]^{\mathcal{T}}$, we show that the complement is open. So, suppose $f(\varepsilon) \neq 1$ or $f(\alpha+\beta) \neq f(\alpha) + f(\beta)$ for some $\alpha, \beta, \alpha+\beta \in \mathcal{T}_0$. If $f(\varepsilon) \neq 1$, then there exists an open set $O \subseteq [0,\infty]$ with $f(\varepsilon) \in O$ and $1 \notin O$ (since $[0,\infty]$ is T_1), and $\pi_\varepsilon^{-1}(O)$ (where π_ε projects f on $f(\varepsilon)$) is an open set contained in the complement of $\mathcal{M}(\mathcal{T}_0)$. Hence, the complement of $\mathcal{M}(\mathcal{T}_0)$ is open. A similar argument can be applied if $f(\alpha+\beta) \neq f(\alpha) + f(\beta)$, and we conclude that $\mathcal{M}(\mathcal{T}_0)$ is closed.

To prove that any intersection of finitely many $\mathcal{M}(\mathcal{T}_0)$'s is nonempty, note that, if $\mathcal{T}_1, \ldots, \mathcal{T}_n$ are finite subsets of $\mathcal{T}$ containg ε, then

$$\bigcap_{i=1}^{n} \mathcal{M}(\mathcal{T}_i) \supseteq \mathcal{M}(\bigcup_{i=1}^{n} \mathcal{T}_i).$$

Since Lemma 3.6 shows that $\mathcal{M}(\bigcup_{i=1}^{n} \mathcal{T}_i) \neq \varnothing$, we therefore find that $\bigcap_{i=1}^{n} \mathcal{M}(\mathcal{T}_i)$ is nonempty, as required.

□

Now, all the hard work is done in proving the theorem, describing the ultimate connection between measure theory and paradoxicality.

Theorem 3.8 (Tarski's Theorem)(AC) *Suppose G is a group acting on a set X, and let $E \subseteq X$. Then E is not G-paradoxical if, and only if, there exists a finitely additive and G-invariant measure μ on $(X, \mathcal{P}(X))$ with $\mu(E) = 1$.*

PROOF: The easy direction was already discussed earlier, so assume that E is not G-paradoxical. Let $\mathcal{S}$ be the type semigroup with respect to G and X and let $\varepsilon = [E \times \{0\}]$. As mentioned earlier, the fact that E is not G-paradoxical means that $\varepsilon = [E \times \{0\}] \neq [E \times \{0,1\}] = 2\varepsilon$. Corollary 3.5 now gives $(n+1)\varepsilon \not\leq n\varepsilon$ for all $n \in \mathbb{N}_0$, and Theorem 3.7 provides a finitely additive measure ν on $\mathcal{S}$ with $\nu(\varepsilon) = 1$ and $\nu(\alpha + \beta) = \nu(\alpha) + \nu(\beta)$ for all $\alpha, \beta \in \mathcal{S}$. Now, the measure $\mu \colon \mathcal{P}(X) \to [0, \infty]$ defined by $\mu(A) = \nu([A \times \{0\}])$ is the desired G-invariant measure.

□

Tarski's theorem shows that no finite groups are paradoxical, since any such group G can be equipped with the measure $\mu(A) = |A|/|G|$. Note in connection with this the introductory remarks concerning infinity and paradoxes.

As another example of an application of Tarski's theorem, one could show that the group $(\mathbb{Z}, +)$ is not paradoxical. This would, of course, be done by constructing a finitely additive measure π with $\pi(\mathbb{Z}) = 1$ and $\pi(E) = \pi(z + E)$ for all $E \subseteq \mathbb{Z}$ and $z \in \mathbb{Z}$. The strategy is motivated by the finite case, as one would attempt to approximate a measure in the form '$E \mapsto |E|/|\mathbb{Z}|$'. Similarly, one could show that $\mathbb{Z}_2 * \mathbb{Z}_2$ is not paradoxical — a rather remarkable result, as we already know that $\mathbb{Z} * \mathbb{Z}$ and $\mathbb{Z}_2 * \mathbb{Z}_3$ are paradoxical.

Our final application of Tarski's theorem is to the Banach–Tarski paradox. In its strong form, the Banach–Tarski paradox combined with Tarski's theorem implies the following amazing result:

Theorem 3.9 (AC) *There exists no finitely additive and* O_3*-invariant measure on* $\mathbb{R}^3$ *that can normalize a bounded set with a nonempty interior.*

Bibliography

1. Stig Alling: *Banach–Tarskis Paradoks [The Banach–Tarski Paradox]* (in Danish), unpublished.

2. Jacob von Bornemann Hjelmborg: *Banach–Tarskis sætning og mængdelære [Banach–Tarski's theorem and set theory]* (in Danish), University of Odense, Odense 1993.

3. Laszlo Lovasz: *Combinatorial Problems and Exercises*, Akadémiai Kiadó, Budapest, and North-Holland Publishing Company, Amsterdam 1979.

4. W.R. Scott: *Group Theory*, Dover Publications, 1987.

5. Stan Wagon: *The Banach–Tarski Paradox*, Cambridge University Press, Cambridge 1985.