Amenability and paradoxicality of groups

Bachelor's thesis

author:
Lennert den Besten

supervisor:
Martijn Caspers

June 2017

Abstract

A group is called paradoxical if it can be decomposed into a finite number of subsets that, after applying a left multiplication, can be reassembled to form two copies of the original group. A group is called amenable if there exists a finitely additive measure on its powerset that normalizes the group and is invariant under left multiplication. The largest part of this thesis is devoted to proving Tarski's theorem, which says that a group is paradoxical if and only if it is not amenable. One direction of this equivalence follows fairly easily after stating the proper definitions, the other requires a few chapters of work, one of which shall be devoted entirely to proving a result in graph theory called König's theorem. The non graph theoretic part of the proof will, in the beginning, mostly consist of generalization and writing down the right definitions. The latter part of the proof will be spent proving properties of a particular semigroup, whose addition is strongly related to the notion of paradoxicality. After we have proved Tarski's theorem we will be looking at left invariant means on locally compact groups. The existence of such a mean generalizes the notion of amenability for locally compact groups. We will then spend the last part of this thesis proving that the existence of a left invariant mean on a locally compact group implies the existence of a left invariant mean on all of its closed subgroups. Which is a result that can be easily shown to imply that any subgroup of an amenable group is amenable too.

Contents

1 Amenability and paradoxicality

In this short first chapter our aim is to introduce the notions of paradoxicality and amenability, and to prove one direction of the equivalence that forms the main problem of this thesis, that is, a group is paradoxical if and only if it is not amenable. In order to do this we will first introduce a few accompanying definitions, and along with this make some relevant remarks. After we have done this, we shall give some examples of paradoxical and amenable groups.

Definition 1.1. *Let X be a set and $\mathcal{A} \subseteq \mathcal{P}(X)$, we call $\mathcal{A}$ an* **algebra** *on X if it has the following properties:*

(a) $X \in \mathcal{A}$

(b) If $A \in \mathcal{A}$, then $A^c \in \mathcal{A}$

(c) If $A, B \in \mathcal{A}$, then $A \cup B \in \mathcal{A}$

Remark 1.2. *The powerset of any set is an algebra. In fact, this will be the only algebra we will be looking at in this thesis, the above definition is given mostly for reasons of completeness of the following definition of finitely additive measures.*

Definition 1.3. *Let $\mathcal{A}$ be an algebra on a set X, we call a function $\mu : \mathcal{A} \to [0, \infty]$ a* **finitely additive measure** *if it has the following properties:*

(a) $\mu(\emptyset) = 0$

(b) If $A, B \in \mathcal{A}$ disjoint, then $\mu(A \cup B) = \mu(A) + \mu(B)$

Remark 1.4. *Above, we defined the codomain of a finitely additive measure to be $[0, \infty]$. For now, just assume it is the space $[0, \infty)$ as we know it, together with a new element that we denote with ∞, together with the convention that $\infty + \infty = \infty$. These are the only properties we will need to know about this space right now. We will study this space in a bit more detail in a later chapter, in order to not fragment definitions and their related theorems too much we shall hold off on giving the precise definition for now.*

Definition 1.5. *Let G be a group acting on a set X and suppose $E \subseteq X$. Then E is called* G-**negligible** *if every finitely additive, left invariant (i.e. $\mu(gA) = \mu(A)$ for all $g \in G$ and $A \in \mathcal{P}(X)$) measure μ on $\mathcal{P}(X)$ is either infinite on E or vanishes on E (i.e. either $\mu(E) = \infty$ or $\mu(E) = 0$).*

Definition 1.6. *Let G be a group, we call G* **amenable** *if there exists a finitely additive, left-invariant measure μ on $\mathcal{P}(G)$ that normalizes G (i.e. $\mu(G) = 1$).*

Remark 1.7. *Consider a group G as acting on itself (as a set) through left multiplication, it is then easy to see that G (as a set) is G-negligible if and only if G (as a group) is not amenable.*

Definition 1.8. *Let G be a group acting on a set X and suppose $E \subseteq X$. Then E is called G-***paradoxical** *if there exist subsets $A_1, ..., A_n, B_1, ..., B_m$ that partition E and elements $g_1, ..., g_n, h_1, ..., h_m$ in G such that both $g_1(A_1), ..., g_n(A_n)$ and $h_1(B_1), ..., h_m(B_m)$ are partitions of E. Furthermore, we shall call a group G paradoxical if in the above description we have that $G = X = E$ and the action is left multiplication.*

Theorem 1.9. *Let G be a group acting on a set X. A paradoxical subset $E \subseteq X$ is G-negligible.*

Proof. Use μ to denote the measure on G and suppose that $\mu(E) \neq \infty$. Let E being paradoxical be witnessed by the subsets $A_1, ..., A_n, B_1, ..., B_m$ and the elements $g_1, ..., g_n, h_1, ..., h_m$. Now, using the disjointness of the subsets that form the partition and property (b) from definition 1.3 we get,

$$\begin{aligned}
\mu(E) &= \mu\left(\left(\bigcup A_i\right) \cup \left(\bigcup B_j\right)\right) \\
&= \sum \mu(A_i) + \sum \mu(B_j) \\
&= \sum \mu(g_i(A_i)) + \sum \mu(h_j(B_j)) \\
&= \mu\left(\bigcup g_i(A_i)\right) + \mu\left(\bigcup h_j(B_j)\right) \\
&= \mu(E) + \mu(E)
\end{aligned}$$

which implies that either $\mu(E) = \infty$ or $\mu(E) = 0$. □

Theorem 1.10. *A paradoxical group is not amenable*

Proof. This is a straightforward consequence of theorem 1.9 combined with remark 1.7 . □

We shall now give a few examples of groups that can fairly easily be identified as either paradoxical or amenable. The first group, or rather set of groups we want to consider are the finite groups. If G is a finite group we can give a specific measure that proves G's amenability, for $A \in \mathcal{P}(G)$ let $\mu(A) = |A|/|G|$, then μ can be easily seen to possess the desired properties. That G is not paradoxical can also be easily seen because paradoxicality would imply that $|G| = 2|G|$, which is not true for finite cardinalities.

In general it is not very simple to directly point out an appropriate measure for an amenable group. There are other equivalent definitions of amenability that are a bit easier to work with. One of them says that a group G is amenable if there exists an infinite sequence $A_1, A_2, ...$ of finite subsets of G such that for every $g \in G$ there exists an $n > 0$ such that $g \in A_i$ for all $i \geq n$, and with the property that

$$\lim_{i \to \infty} \frac{|gA_i \Delta A_i|}{|A_i|} = 0 \qquad \text{for all } g \in G,$$

where Δ denotes the symmetric difference. Using this definition we can quickly see that the group $\mathbb{Z}$ is amenable, just take $A_i = \{x \in \mathbb{Z} \mid |x| \leq i\}$, that the above limit then vanishes is not hard to see. Something that is hard, is trying to explicitly write down an appropriate measure for $\mathbb{Z}$ to show its amenability, even though $\mathbb{Z}$ a group that is not considered very complicated. Note that we will not be proving the equivalence of this alternative definition of amenability (which is called the Fölner condition) or use it anywhere outside of this example.

An example of a paradoxical group is the free group of rank 2, which we will denote with $\mathbb{F}_2$ and whose two generating elements we shall call a and b. A paradoxical decomposition can be made as follows, let

$$\begin{aligned}
A_1 &= \{\text{all words that start with } a^n \text{ for } n \geq 1\} \\
A_2 &= \{\text{all words that start with } a^n \text{ for } n \leq 1\} \\
B_1 &= \{\text{all words that start with } b^n \text{ for } n \geq 1\} - \{b, b^2, ...\} \\
B_2 &= \{\text{all words that start with } b^n \text{ for } n \leq 1\} \cup \{e, b, b^2, ...\}.
\end{aligned}$$

and note that A_1, A_2, B_1, B_2 form a partition of $\mathbb{F}_2$. But also note that both $a^{-1}A_1, eA_2$ and $b^{-1}B_1, eB_2$ form partitions of $\mathbb{F}_2$, this shows paradoxicality.

2 König's theorem, a result in graph theory

We shall devote this entire chapter to proving a result in graph theory called König's theorem. This result will play an important role in the proof of a theorem in the next chapter. First, though, we will have to formulate a dozen or so definitions and prove two other results. We will not presuppose any graph-theoretic knowledge and begin with the definition of a graph. However, we shall not be illustrating the definitions through examples for the sake of brevity. If any definition is unclear we refer to [1], the definitions are fairly analogous to the ones we use here and an abundance of examples is provided.

A **graph** is made up of two sets, first a set V, whose elements we shall call **vertices**, together with a collection E containing unordered pairs of V, we call these pairs in E **edges**. We want to allow our graphs to have multiple edges for the same pair of vertices, so E is allowed to contain the same element multiple times. If e is an edge we shall call the vertices e contains its **ends**. The **degree** of a vertex v is the number of edges that contain v. If (V, E) is a graph, we shall call any graph (V', E'), with the property that $V' \subseteq V$ and $E' \subseteq E$, a **subgraph** of (V, E).

Let (V, E) be a graph, we call a subset M of E a **matching** if all edges in M have two distinct ends and no two edges share an end. If $V' \subseteq V$ is a set of vertices that are all an end of a particular matching M, we call V' and its elements M**-saturated**. We shall furthermore call a matching M **perfect** if V is M-saturated and call M **maximum** if there exists no matching M' such that $|M'| > |M|$.

We define a **path** to be a finite sequence of edges with two properties. Firstly that each edge in the sequence shares one of its ends with its predecessor (when it has one) and shares its other end with its successor (when it has one). Secondly, all the edges of a path must be distinct and all the shared ends must also be distinct. Note, that within this definition we don't bother with specifying the direction of single-edge paths, for our purposes just knowing the length of a particular path suffices.

Let M be a matching in a graph, we call a path M**-alternating** if its edges are in M and not in M alternately. If the first and last edge of an M-alternating path are both not edges in M and the **start** and **endpoints** (the ends of the first and last edges in the path that are not shared with their successor and predecessor, respectively) are not M-saturated we call it an M**-augmenting** path.

We call two vertices **connected** if there exists a path with them as start and endpoints. Note that connectedness is an equivalence relation, we call the subgraphs generated by the equivalence classes (that is, the graph with the equivalence class as set of vertices, and as set of edges the set containing precisely the edges that have an end in the equivalence class) the **components** of a graph.

We are now ready to state and prove the first theorem of this chapter, after which we will have to introduce a few new definitions before proving the lemma and König's theorem.

Theorem 2.1. *(Berge) Let M be a matching in some finite graph, then M is maximum if and only if there is exists no M-augmenting path.*

Proof. Let's first suppose that M is maximum and that there exists an M-augmenting path $e_1, ..., e_n$. First note that the set of edges $M' = (M - \{e_i \mid i \text{ is even}\}) \cup \{e_i \mid i \text{ is odd}\}$ is a matching because the start and endpoints of the path are not M-saturated. Then note that since e_1 and e_n are both not edges in M and the path is M-alternating, n has to be odd. Therefore there is exactly one more e_i with odd index than with even index. So our new path M' is exactly one edge longer than M, this is however in contradiction with the premise that M is maximum.

Conversely, suppose that there exists no M-augmenting path and that M is not maximum. Let M' be a matching that is maximum, this tells us that $|M'| > |M|$. Now let $M'' = M' \Delta M$ and let V be the set containing the ends of all the edges in M''. Each vertex in V is the end of at least one edge in M'' and at most two edges in M'' which are not both in M or M'. Therefore, the set of edges of every component of M'' form a path whose edges are in M and M' alternately. Notice that since $|M'| > |M|$ there must be such a component C that has more edges in M' than edges in M. The path induced by C has to have an odd number of vertices (this follows from the alternating property), therefore the start and endpoints of this path are distinct. The start and endpoints can also not be M-saturated, for if this was the case there would exist an edge adjacent to a vertex in C but not included in it. We can conclude that this path is M-augmenting, this however is in contradiction with our premise. □

Definition 2.2. *We call a graph* **bipartite** *if its set of vertices can be split into two parts such that every edge has one end in either part, and we call a graph k-**regular** if all its vertices have degree $k > 0$.*

Definition 2.3. *Let V be a set containing some vertices of a graph. We call the set containing exactly the vertices that are adjacent to a vertex in V (i.e. there exists an edge with these vertices as ends) the* **neighbor set** *of V, we will denote this set with $N(V)$.*

Lemma 2.4. *A finite, bipartite graph with bipartition A, B has a matching that saturates A if and only if $|N(S)| \geq |S|$ for all $S \subseteq A$.*

Proof. First let M be a matching that saturates A and let $S \subseteq A$. Then M also saturates S, in other words every vertex in S is the end of some edge e in M . Such an e will have its other end in B (and therefore not in S) because A and B are a bipartition. This tells us that the neighbor set of S contains all such e and is at least as big as S itself.

Conversely, suppose that there exists no matching that saturates A but that $|N(S)| \geq |S|$ for all $S \subseteq A$. Let M be a maximum matching, since M does not saturate A there exists an $a \in A$ such that a is not M-saturated. Denote with C the set of all vertices that are connected to a by an M-alternating path. Because of theorem 2.1 no such path will be M-augmenting. So, all vertices in C except a are M-saturated. Let $U = A \cap C$ and $V = B \cap C$, then note that every $u \in U - \{a\}$ is matched to a unique $v \in V$ by an edge in M. This tells us that U is a larger set than V and also that $V \subseteq N(U)$. We actually claim that the reverse inclusion also holds. To prove this we have to show the existence of an M-alternating path from a to every vertex in $N(U)$. Suppose $e_1, ..., e_n$ is an M-alternating path from a with endpoint u in U and let e be an edge between u and a vertex u' in $N(U)$. We know from before that e_n has to be an edge in M. Now, if e is not in M, the sequence $e_1, ..., e_n, e$ is an M-alternating path that connects a with u'. If e is in M then e has to equal to e_n because M is a matching, in this case $e_1, ..., e_{n-1}$ is an

M-alternating path connecting a with u'. This proves our claim. Note that this claim actually implies that $N(U)$ is equal to V and is therefore smaller than U. This is a contradiction because $U \subseteq A$ implies that $|N(U)| \geq |U|$.

□

Theorem 2.5. *(König's Theorem) A k-regular bipartite graph has perfect matching.*

Proof. We will first prove the theorem for finite graphs. So, let (V, E) be a finite k-regular graph with bipartition A, B. Then, using regularity we get $k|A| = |E| = k|B|$ which implies that A and B are equal in size. Let $S \subseteq A$, denote with E' the set of edges that have an end in S and denote with E'' the set of edges that have an end in $N(S)$, note here that $E' \subseteq E''$. Using regularity again tells us that $k|N(S)| = |E''| \geq |E'| = k|S|$ which implies that $|N(S)| \geq |S|$. This, together with the assumption that (V, E) is finite and bipartite means that we can apply lemma 2.4. It tells us that there exists a matching that saturates A, because $|A| = |B|$ this matching is perfect. This concludes the finite case.

For our infinite case let us first note that an infinite k-regular graph only has countable components, because for any vertex a path of length n can at most reach k^n other vertices. Also note that to actually obtain an enumeration of every component of a graph that consists of uncountable many components, the axiom of choice is needed. We are looking at the components of graphs because for a particular graph our conclusion is true if and only if it is true in every component of this graph. Combining this with our earlier remark about countability means that the theorem is true for all graphs if it is true for the countable ones.

So, let (V, E) be a countable graph and let $e_1, e_2, ...$ enumerate E. We claim that for every segment $e_1, e_2, ..., e_n$ there exists a graph (V'_n, E'_n) (not necessarily a subgraph of (V, E)) that is k-regular, bipartite and such that E'_n contains all $e_1, e_2, ..., e_n$. This claim is true because the graph generated by $e_1, e_2, ..., e_n$ and its adjacent vertices is finite, bipartite and every vertex has at most degree k. Any such graph can be extended to being finite, k-regular and bipartite by first adding vertices to one part of the bipartition (if the parts aren't equal already) and by then adding edges to bring every vertex degree up to k. This last part is possible because bipartiteness implies that the two parts have an equal amount of adjacent edges, this combined with the fact that our graphs are allowed to have multiple edges means that we can just keep adding edges between the vertices in the two parts that have degree lower than k until they do have degree k. Now, using the finite version of the theorem tells us that all such (V'_n, E'_n) have perfect matching. The idea is to use these perfect matchings to inductively create a perfect matching for (V, E).

To help us write down this perfect matching we will use finite sequences of 0's and 1's. We call such a sequence s_m of length m good if there exists a graph (V'_m, E'_m) with properties exactly as before, such that $s_m(i) = 1$ if and only if e_i is an edge in a perfect matching of (V'_m, E'_m). Note that any initial segment of a good sequence is also good. Also, because of our earlier claim there exists at least one good sequence of every length. We can now inductively construct an infinite sequence s so that every initial segment of s is good. Choose $s(l)$ to be 0 or 1 such that $s(1)...s(l)$ is good and has infinitely many good extensions, if this holds for both 0 and 1 let $s(l) = 0$.

It is not entirely obvious that such a choice can always be made. An illustrative way to prove that this is possible is by using induction. So suppose that $l = 1$, because there are infinitely many good sequences there have to be infinitely many good sequences that all start with 0 or all start with 1. Without loss of generality, assume that this is true for 0 and let $s(1) = 0$. The first element in a good sequence is an initial segment and is therefore also a good sequence, so we have a good sequence of length one. We have assumed earlier that there are infinitely many good sequences that start with 0, all such sequences that are not single-length (of which there must be an infinite amount) are extensions of the sequence $s(1)$. This concludes the first induction step. Now suppose that $s(1)...s(l-1)$ is good and has infinitely many good extensions, we can apply more or less the same reasoning as above. Again, because there are infinitely many good extensions there must be infinitely many whose l-th element is 0 or 1. Let's say, without loss of generality, that this holds for 0, and let $s(l) = 0$. Then $s(1)...s(l)$ is good because it is an initial segment of a good sequence. Moreover, it is an initial segment of all the good extensions of $s(1)...s(l-1)$ that are more than l elements long and whose l-th element is 0. Since there is just one such extension that is exactly l elements long, our earlier assumption implies that there are infinitely many good extensions of $s(1)...s(l)$. This means that we are done.

Now that we have our infinitely long good sequence s we can use it to construct a perfect matching. Let $M = \{e_i \mid s(i) = 1\}$, that this is a perfect matching can be seen as follows. Let v be a vertex in V and let l be so large that all k edges that are adjacent to v are elements in the sequence $e_1, ..., e_l$. Because the sequence $s(1)...s(l)$ is good, v is matched by M.

□

3 Equidecomposability and the type semigroup

In this chapter we will begin with generalizing the concept of paradoxicality. Intuitively we would like there to be a way to add sets such that a set is if paradoxical if and only if it doesn't change when its added onto itself. It turns out that such an addition exists when taking a countably infinite amount of copies of a set and considering the notion of equidecomposability on this new space. After making all of this precise we will be looking at some of the properties of this addition. Some properties are very straightforward while others will require a fair amount of work. Let's start with discussing the notion of equidecomposability.

Definition 3.1. *Let G be a group acting on a set X and suppose $A, B \subseteq X$. Then A and B are called G-**equidecomposable** if there exist $A_1, ..., A_n$ partition of A, $B_1, ..., B_n$ partition of B and $g_1, ..., g_n$ elements in G such that $g_i(A_i) = B_i$ for each $i \leq n$.*

We use the notation $A \sim B$ to denote that two sets are equidecomposable (we will sometimes write $\sim_G$ to specify which group we are talking about). Note that $\sim$ is an equivalence relation, we will check the transitive property right now (reflexivity and symmetry are quite straightforward). Suppose that $A \sim B$ using the notations as in the above definition and suppose that $B \sim C$ witnessed by $B'_1, ..., B'_m$ partition of B, $C_1, ..., C_m$ partition of C and $h_1, ..., h_m$ elements in G. We want to prove that $A \sim C$, to do this we will write down specific partitions elements that witness the equidecomposability relation. First we define a new partition of A by introducing an upper index on our A_i's. Define $A_i^j = A_i \cap g_i^{-1} B'_j$ for $i \leq n$ and $j \leq m$. Note that these sets are disjoint because all A_i's are disjoint (by definition) and all $g_i^{-1} B'_j$'s are disjoint over varying j (all B'_j being disjoint by definition implies that all $g_i^{-1} B'_j$'s are also disjoint when i is kept fixed). That these A_i^j form A under union (and thus partition A) can be seen as follows

$$\bigcup_{i \leq n} \bigcup_{j \leq m} A_i^j = \bigcup_{i \leq n} \bigcup_{j \leq m} (A_i \cap g_i^{-1} B'_j) = \bigcup_{i \leq n} A_i \cap (g_i^{-1} \bigcup_{j \leq m} B'_j) = \bigcup_{i \leq n} (A_i \cap g_i^{-1} B) = \bigcup_{i \leq n} A_i = A.$$

To get an appropriate sequence of group elements we apply the same idea. Let $x_i^j = h_j g_i$ for $i \leq n$ and $j \leq m$, we then get

$$\begin{aligned}
\bigcup_{j \leq m} \bigcup_{i \leq n} x_i^j A_i^j &= \bigcup_{j \leq m} \bigcup_{i \leq n} h_j g_i (A_i \cap g_i^{-1} B'_j) \\
&= \bigcup_{j \leq m} h_j \bigcup_{i \leq n} g_i (A_i \cap g_i^{-1} B'_j) \\
&= \bigcup_{j \leq m} h_j \bigcup_{i \leq n} (g_i A_i \cap B'_j) \\
&= \bigcup_{j \leq m} h_j ((\bigcup_{i \leq n} g_i A_i) \cap B'_j) \\
&= \bigcup_{j \leq m} h_j (B \cap B'_j) \\
&= \bigcup_{j \leq m} h_j B'_j \\
&= B.
\end{aligned}$$

Now, what is left to show to prove equidecomposability of A and C, is that all $x_i^j A_i^j$'s are disjoint. For this, note that $x_i^j A_i^j = h_j(g_i A_i \cap B_j') = (h_j g_i A_i \cap h_j B_j') = (h_j B_i \cap C_j)$ and that these sets are disjoint because of the same reasoning we used to show that all A_i^j's are disjoint.

Definition 3.2. *Let G be a group acting on a set X. We introduce the following notation: let $X^* = X \times \mathbb{N}$ and $G^* = \{(g, \pi) \mid g \in G \text{ and } \pi \text{ is a permutation of } \mathbb{N}\}$. Note that G^* is a group that acts naturally on X^* by the following action $(g, \pi)(x, n) = (g(x), \pi(n))$.*
*Let $A \subseteq X^*$, we call the $n \in \mathbb{N}$ for which $A \cap (X \times \{n\}) \neq \emptyset$ the **levels** of A. If A has finitely many levels, we call it* **bounded**.

The idea behind the above definition is that it allows us to copy and break up subsets of X and spread them over different levels. When considering the concept of G^*-equidecomposability of subsets of X^*, one can note that it gives rise to a fairly concise notion of paradoxicality of subsets of X. A set $A \subseteq X$ is G-paradoxical if and only if $A \times \{n\} \sim_{G^*} A \times \{k, m\}$, with $n, m, k \in \mathbb{N}$ (note that the choice of natural numbers doesn't matter as $A \times \{n\} \sim_{G^*} A \times \{k\}$ for any $A \subseteq X$ and $n, k \in \mathbb{N}$). We can actually define an addition on the equivalence class (with respect to $\sim_{G^*}$) of subsets of X^* that have finitely many levels. When adding two sets we can just shift the levels of one of them so that their levels don't overlap. This makes sense because just shifting the levels of a bounded subset of X^* will not change its equivalence class. Let us make this precise and show that all of this is actually true and well-defined.

Definition 3.3. *Use general notation as above. Suppose A is bounded subset of X^*, we denote the equivalence class (with respect to $\sim_{G^*}$) of A with $[A]$, we call this the* **type** *of A. Furthermore, denote the set containing all types as $\mathcal{S}$. We can define addition on $\mathcal{S}$ as follows: suppose that $[A], [B] \in \mathcal{S}$, then let $[A] + [B] = [A \cup B']$ where B' equals $(e, \pi)(B)$, where π is any permutation of $\mathbb{N}$ that sends levels that A and B share to a level that none of them have and leaves the other levels of B fixed.*

We want to show that this addition is well defined and that $(\mathcal{S}, +)$ in fact forms a commutative **semigroup** (i.e. '+' is associative) with identity. We shall refer to $(\mathcal{S}, +)$ as the **type semigroup**.

Let's start with well-definedness, for this we have to show that $[A] + [B]$ is independent of choice of representatives of $[A]$ and $[B]$. So, suppose that $A_1 \sim A_2$ (both in $[A]$), witnessed by $A_1^1, ..., A_1^n$ partition of A_1 and $A_2^1, ..., A_2^n$ partition of A_2 and $g_1, ..., g_n$ elements in G^*. Because A_1 and A_2 are both bounded, there exists a shift B' of B that doesn't share any levels with both A_1 and A_2. Then note that $A_1^1, ..., A_1^n, B'$ partitions $A \cup B'$ and that $g_1(A_1^1), ..., g_n(A_1^n), e(B') = A_2^1, ..., A_2^n, B'$ which implies that $A_1 \cup B' \sim A_2 \cup B'$, or equivalently that $[A_1] + [B] = [A_1 \cup B'] = [A_2 \cup B'] = [A_2] + [B]$. Now suppose that $B_1 \sim B_2$ (both in $[B]$), again we can argue that there exists a shift B' of B_1 that doesn't share levels with both A and B_2 because the boundedness property. We can apply the same reasoning as before to see that $[A] + [B_1] = [A] + [B_2]$. This shows that '+' is well-defined.

For commutativity suppose that A and B are bounded and let B' be an suitable shift of B. Then note that since $B \sim B'$, we have $[A] + [B] = [A \cup B'] = [B' \cup A] = [B'] + [A] = [B] + [A]$. Now suppose that C is also bounded, then there exists a shift C' that is doesn't share any levels with either A or B' because both these sets have finite levels. Then $([A] + [B]) + [C] = [A \cup B'] + [C] = [A \cup B' \cup C'] = [A] + [B' \cup C'] = [A] + ([B'] + [C']) = [A] + ([B] + [C])$ which shows associativity. Lastly note that the equivalence class of empty set acts as identity element since $[A] + [\emptyset] = [A \cup \emptyset] = [A]$.

Remark 3.4. *Again, use general notation as above. A subset $E \subseteq X$ is G-paradoxical if and only if $[E \times \{0\}] = 2[E \times \{0\}]$*

Proof. First suppose that E is G-paradoxical, witnessed by the partition $A_1, .., A_n, B_1, ..., B_m$ of E and the elements $g_1, ..., g_n, h_1, ..., h_m$ in G. Then, let $A_i^* = A_i \times \{0\}$ for $i \leq n$ and $B_j^* = B_j \times \{0\}$ for $j \leq m$, these subsets of X^* form a partition of $E \times \{0\}$. Likewise, let $g_i^* = (g_i, \pi_1)$, for $i \leq n$ where π_1 is the identity permutation of $\mathcal{N}$ and let $h_j^* = (h_j, \pi_2)$ for $j \leq m$ where π_2 is a permutation of $\mathcal{N}$ that sends 0 to 1. These sets and elements marked with a star then witness that $E \times \{0\} \sim E \times \{0, 1\}$, note that the right hand side of this equivalence is a representative of $2[E \times \{0\}]$. This tells us that $[E \times \{0\}] = 2[E \times \{0\}]$.

Now suppose that $[E \times \{0\}] = 2[E \times \{0\}]$, witnessed by the partition $A_1 \times \{0\}, .., A_n \times \{0\}$ of E and elements $(g_1, \pi_1), ..., (g_n, \pi_n)$. The permutations π_i will send 0 either to itself or to 1, we can sort the A_i's based on this fact. Let $I_0 = \{i \leq n \mid \pi_i(0) = 0\}$ and $I_1 = \{i \leq n \mid \pi_i(0) = 1\}$. Then the all of the subsets $A_1, ..., A_n$ and $\{g_i(A_i)\}_{i \in I_1}$ and $\{g_i(A_i)\}_{i \in I_2}$ form partitions of E which proves G-paradoxicality of E. □

Above, we only shortly mentioned the definition of a semigroup. We want to say a little bit more about semigroups however. Namely, any commutative semigroup with identity admits a natural way of multiplying elements with natural numbers and also a natural way of ordering elements.

Let $(\mathcal{T}, +)$ be a commutative semigroup with identity 0 and let $n \in \mathbb{N}$. For an element $\alpha \in \mathcal{T}$ we define $n\alpha = 0$ when $n = 0$ and for $n > 0$ we let $n\alpha = \alpha + ... + \alpha$, where we sum n times on the right hand side. It can be easily seen that associative (i.e. $n(m\alpha) = (nm)\alpha$) and distributive (i.e. $n(\alpha + \beta) = n\alpha + n\beta$) laws are satisfied.

The semigroup can be ordered as follows. Let $\alpha, \beta \in \mathcal{T}$, we say that $\alpha \leq \beta$ if there is a $\gamma \in \mathcal{T}$ such that $\alpha + \gamma = \beta$. With this definition we can, like before, quickly see some properties that this order $\leq$ has. For instance $n\alpha \leq n\beta$ if $\alpha \leq \beta$, $n\alpha \leq m\beta$ if $n \leq m$ and $\alpha + \gamma \leq \beta + \gamma$ if $\alpha \leq \beta$. An important property that the type semigroup has with regard to multiplication has is that $n\alpha = n\beta$ implies that $\alpha = \beta$. That this property holds, however, is not as easily seen as the ones we just stated above. We shall discuss this property in our next theorem. Before we state this theorem however, we will discuss two short facts on equidecomposability that we shall use in its proof.

Lemma 3.5. *Let G be a group acting on a set X and suppose $A, B, C, D \subseteq X$, then the following two statements are true.*

(a) If $A \sim B$, witnessed by $A_1, ..., A_n$, $B_1, ..., B_n$ and $g_1, ..., g_n$, then there is a bijection $f : A \to B$ induced by this decomposition that satisfies the property that if $C \subseteq A$, then $C \sim f(C)$.

(b) If $A \sim B$ and $C \sim D$ such that $A \cap C = \emptyset = B \cap D$, then $A \cup C \sim B \cup D$.

Proof. Firstly, (a) can be seen to be true because we can build such an f from all the group actions $g_i : A_i \to B_i$. Remember here that all the A_i's and B_i's form partitions. Also note that $f^{-1} : A \to B$ can be seen as being induced by the same sets and elements $g_1^{-1}, ..., g_n^{-1}$ and therefore holds the same property regarding subsets of B.

The second fact has been argued more or less while discussing the well-definedness of the addition of $(\mathcal{S}, +)$. □

Theorem 3.6. *(Cancellation law) Suppose $\alpha, \beta \in \mathcal{S}$, then $n\alpha = n\beta$ implies that $\alpha = \beta$*

Proof. Let's first interpret some of the notation. Writing $n\alpha = n\beta$ means that there exist sets $A_1, ..., A_n$ disjoint, bounded subsets of X^* such that $[A_i] = \alpha$ for all $i \leq n$, and sets $B_1, ..., B_n$ with identical properties with respect to β, such that $[A_1 \cup ... \cup A_n] = [B_1 \cup ... \cup B_n]$. Let $A = A_1 \cup ... \cup A_n$ and $B = B_1 \cup ... \cup B_n$, then $A \sim B$.

For each equidecomposable relation we just mentioned, there exists an induced bijection like described earlier. Let $\chi : A \to B$ and $\phi_i : A_1 \to A_i$, $\psi_i : B_1 \to B_i$ be such bijections. Using these ϕ_i and ψ_i we can construct new partitions of A and B. For $a_1 \in A_1$ and $b_1 \in B_1$, define $\overline{a_1} = \{\phi_1(a_1), ..., \phi_n(a_1)\}$ and $\overline{b_1} = \{\psi_1(b_1), ..., \psi_n(b_1)\}$. Then the collection $\{\overline{a_1} \mid a_1 \in A_1\}$ partitions A and likewise $\{\overline{b_1} \mid b_1 \in B_1\}$ partitions B.

We have now reached the point where we want to apply the graph-theoretic knowledge we acquired in chapter 2 . We can build a bipartite graph as follows. Start with A_1, B_1 as the bipartition of the set of vertices. For every $a_1 \in A_1$ and each $i \leq n$ create an edge that connects a_1 to the $b_1 \in B_1$ for which there exists a $j \leq n$ such that $\chi\phi_i(a_1) = \psi_j(b_1)$, this way each a_1 has degree n. We actually claim that each b_1 also has degree n. This is true because the edges of this graph are induced by the bijection χ. More precisely, each $a \in A$ induces exactly one edge in our graph (because $\{\overline{a_1} \mid a_1 \in A_1\}$ partitions A) and therefore each $b \in B$ does the same. For each $b_1 \in B_1$ there are n distinct b such that $\psi_j(b_1) = b$ for a $j \leq n$. This actually means that our graph is n-regular as well as being bipartite. Therefore, König's theorem applies and a perfect matching M exists.

Every edge in M uniquely connects an $a_1 \in A_1$ to a unique $b_1 \in B_1$. Recall the way the edges in our graph are defined, for every such edge $\{a_1, b_1\}$ there exist unique $i, j \leq n$ such that $\chi\phi_i(a_1) = \psi_j(b_1)$. Because M is perfect, we can partition A_1 by sorting elements depending on which ϕ_i and ψ_j are used to induce the unique connecting edge in M. We can do this by introducing an upper index on A_1, let $A_1^{ij} = \{a_1 \in A_1 \mid a_1 \text{ is matched by virtue of } \phi_i \text{ and } \psi_j\}$, all these A_1^{ij} will partition A_1. Note that because M is perfect, all $\psi_j^{-1}\chi\phi_i(A_1^{ij})$ will actually partition B_1. Furthermore, χ and all ϕ_i and ψ_j are bijections like we described in (a) in the lemma and therefore the equivalence $A_1^{ij} \sim \psi_j^{-1}\chi\phi_i(A_1^{ij})$ holds for all $i, j \leq n$. Now applying property (b) of the lemma on these partitions tells us that $A_1 \sim B_1$, or equivalently, that $\alpha = \beta$, which is what we intended to prove. □

Another property that the ordering of the type semigroup has is that for two elements α, β we have that $\alpha \leq \beta$ together with $\alpha \geq \beta$ implies that $\alpha = \beta$. Proving this luckily doesn't require an entire chapter of graph theory. We will, however, need to introduce some new notation regarding the equidecomposability equivalence relation before we can prove the relevant theorem. With this new notation, the property is a short corollary of the Banach-Schröder-Bernstein theorem, which we shall state and prove along with a few of its other important consequences.

Definition 3.7. *Let A and B subsets of a set that has a group acting on it. We write $A \preceq B$ if A is equidecomposable with a subset of B.*

Remark 3.8. *Let X, G, X^*, G^* be as in definition 3.2 and suppose that A and B are bounded sets in X^*, then $[A] \leq [B]$ if and only if $A \preceq B$.*

Proof. First, suppose that $[A] \leq [B]$, then there exists a bounded set $C \subseteq X^*$ such that $[A]+[C] = [B]$, then there is a bounded set $C' \subseteq X^*$ that is disjoint from A such that $C \sim C'$ and $A \cup C' \sim B$. Let this equidecomposability be witnessed by the sequences $A_1, ..., A_n, C_1, ..., C_m$ and $B_1, ..., B_{n_m}$, and the elements $g_1^*, ..., g_{n+m}^*$. Define $B' = B_1 \cup ... \cup B_n$, then note that $B' \subseteq B$ and that $A \sim B'$ is witnessed by the subsets $A_1, ..., A_n$ and $B_1, ..., B_n$ and elements $g_1^*, ..., g^*n$. This means that $A \preceq B$.

For the other direction of the statement suppose that $A \preceq B$. This means that there is a $B' \subseteq B$ such that $A \sim B'$, suppose that this $\sim$ is witnessed by subsets $A_1, ..., A_n$ and $B'_1, ..., B'_n$, and elements $g_1^*, ..., g^*n$. Let C be a shift of the bounded set $B - B'$ such that C and A are disjoint, denote the element by which $B - B'$ is shifted by $(e, \pi) \in G^*$. Then $A \cup C \sim B$ is witnessed by the subsets $A_1, ..., A_n, C$ and $B'_1, ..., B'_n, B - B'$, and the elements $g_1^*, ..., g^*n, (e, \pi^{-1})$. This tells us that $[A] + [C] = [A \cup C] = [B]$, which means that $[A] \leq [B]$.

∎

Theorem 3.9. *(Banach-Schröder-Bernstein) Let G be a group that acts on a set X and suppose that $A, B \subseteq X$. Then $A \preceq B$ together with $A \succeq B$ implies $A \sim B$.*

Proof. We shall only use the fact that $\sim$ satisfies the two properties stated in lemma 3.5, the theorem actually holds for any equivalence relation on $\mathcal{P}(X)$ that satisfies these properties. Let $A \preceq B$ and $A \succeq B$ be witnessed by $B \sim A' \subseteq A$ and $A \sim B' \subseteq B$, then let $f : A \to B'$ and $g : A' \to B$ be bijections like described in (a) in the lemma. Now, define $C_0 = A - A'$ and inductively let $C_{n+1} = g^{-1}f(C_n)$, then, let $C = \bigcup_{i=o}^{\infty} C_i$. This way an element $a \in A$ lies in C if and only if $a \notin A'$ or there exists an $c \in C$ such that $a = g^{-1}f(c)$. Taking the negation of this last sentence tells us that an element $a \in A$ lies in $A - C$ if and only if $a \in A'$ and for all $c \in C$ we have $a \neq g^{-1}f(c)$. Applying g on this sentence then says that an element $b \in B$ lies in $g(A - C)$ if and only if there exists an element $a \in A'$ such that $g(a) = b$ (which always exists because g is a bijection) and $b \neq f(c)$ for all $c \in C$. In other words, $g(A - C) = B - F(C)$, then, because g is a bijection like described in (a) in the lemma we have $A - C \sim B - F(C)$. Because f is defined in the same manner as g we also have $C \sim F(C)$. Applying property (b) from the lemma then gives $A = A - C \cup C \sim B - f(C) \cup f(C) = B$, which is the equivalence we wanted to show. ∎

Remark 3.10. *The following definitions of paradoxicality are equivalent to the definition we gave in chapter one.*

Let G be a group acting on a set X and suppose $E \subseteq X$. Then E is called G-paradoxical if there exist subsets $A_1, ..., A_n, B_1, ..., B_m$ of E that are pairwise disjoint and elements $g_1, ..., g_n, h_1, ..., h_m$

in G such that $g_1(A_1), ..., g_n(A_n)$ and $h_1(B_1), ..., h_m(B_m)$ both form partitions of E.

Let G be a group acting on a set X and suppose $E \subseteq X$. Then E is called G-paradoxical if there exist subsets $A_1, ..., A_n, B_1, ..., B_m$ of E that are pairwise disjoint and elements $g_1, ..., g_n, h_1, ..., h_m$ in G such that $g_1(A_1) \cup ... \cup g_n(A_n) = E = h_1(B_1) \cup ... \cup h_m(B_m)$.

Proof. For brevity, we shall refer to the original definition as (a) and the above ones as (b) and (c) respectively. That (a) implies (c) can easily be seen, since the restrictions on the sets in (a) are strictly stronger than the ones in (c).

Now we want to show that (c) implies (b). So, suppose that the statement in (c) is true and use identical notation. We can, without loss of generality, assume that the A_i's and B_j's are nonempty. It is fairly easy to see that we can reduce the A_i's and B_j's such that the $g_i(A_i)$'s and $h_j(B_j)$'s form partitions of E. One way to do this would be by letting $A_i' = A_i - g_i^{-1}(\bigcup_{i<j\leq n} g_j(A_j))$ for $i \leq n$ (this way we delete any element from an A_i that g_i sends to an element that will be covered by a $g_j(A_j)$ with higher index) and applying the same process to get a new sequence of B_j''s from the B_j's. Note here that these A_i''s and B_j''s are still pairwise disjoint.

Lastly, for equivalence of all three definitions we have left to prove that (b) implies (a). Use notation as in (b) and let $A = A_1 \cup ... \cup A_n$ and $B = B_1 \cup ... \cup B_m$, then, looking back at the definition of equidecomposability we have $A \sim E \sim B$. Also, if we can show that $E - A \sim E$, we are done. Since $E - A \subseteq E$ we have that $E - A \preceq E$, on the other hand, we have that $E \sim B \subseteq E - A$ which implies that $E \preceq E - A$, applying Banach-Schröder-Bernstein yields $E - A \sim E$, so we are done. □

Corollary 3.11. *Let $\alpha, \beta \in \mathcal{S}$, then, $\alpha \leq \beta$ together with $\alpha \geq \beta$ implies $\alpha = \beta$.*

Proof. The type semigroup $\mathcal{S}$ is induced by the action of a group on some set X. Because of this, there exist $A, B \subseteq X^*$ such that $[A] = \alpha$ and $[B] = \beta$. This means that $\alpha \leq \beta$ and $\alpha \geq \beta$ implies that $[A] \leq [B]$ and $[A] \geq [B]$, which implies that $A \preceq B$ and $A \succeq B$ according to remark 3.8 . Finally, applying Banach-Schröder-Bernstein gives $A \sim B$, which implies $\alpha = [A] = [B] = \beta$. □

Corollary 3.12. *Let $\alpha \in \mathcal{S}$ and $n \in \mathbb{N}$, then, $(n+1)\alpha \leq n\alpha$ implies that $\alpha = 2\alpha$*

Proof. By applying the inequality $(n+1)\alpha \leq n\alpha$ on its own left hand side n times yields $2n\alpha \leq n\alpha$. The reverse inequality $2n\alpha \geq n\alpha$ also holds since $n\alpha + n\alpha = 2n\alpha$. Because of the previous corollary this combines into $2n\alpha = n\alpha$, or $n(2\alpha) = n\alpha$, then applying the cancellation law gives $2\alpha = \alpha$, as desired. □

4 Tarski's theorem

We are now almost ready to state and prove the main theorem of this thesis. Before we do this, though, in order to slightly shorten its still fairly long proof, we shall state and prove some results regarding topology, and finitely additive functions and measures.

We shall be defining **finitely additive** functions on a semigroup $\mathcal{T}$, this with the aim to convert such a function into a measure on the type semigroup. With finite additivity of a function $f : \mathcal{T} \to [0, \infty]$ on a subset $\mathcal{T}' \subseteq \mathcal{T}$, we mean that, if α, β and $\alpha + \beta$ all lie in $\mathcal{T}'$, then f has the property that $f(\alpha + \beta) = f(\alpha) + f(\beta)$. A small fact we want to note is that if $\mathcal{T} = \mathcal{T}'$, then such a function f is order-preserving (i.e. $\alpha \leq \beta$ implies that $f(\alpha) \leq f(\beta)$. This is quite straightforward since $\alpha \leq \beta$ implies that there exists a $\gamma \in \mathcal{T}$ such that $\alpha + \gamma = \beta$, which, after applying f tells us that $f(\alpha) + f(\gamma) = f(\beta)$, which means that $f(\alpha) \leq f(\beta)$ because f only takes on values greater or equal than 0.

Next, recalling our definition of (finitely additive) measures, note that the we defined the codomain of a measure to be the space $[0, \infty]$. For the proof of the next theorem we need to know some of the topological properties of this space. We shall discuss the relevant properties shortly. Before we do this however, we will clarify how exactly $[0, \infty]$ is defined and what topology is endowed upon it.

The set $[0, \infty]$ is defined as the set obtained by the taking set of non-negative reals $[0, \infty)$ and adding a new element denoted by ∞. We let $[0, \infty]$ inherit the ordering of $[0, \infty)$ on the elements they share and order ∞ by letting $x < \infty$ for all $x \in [0, \infty)$. The topology endowed on this space is generated by the base that contains all sets that are open in $[0, \infty)$ and all intervals $(a, \infty]$ where $a \in [0, \infty)$. Note that by this definition any set not containing ∞ is open in $[0, \infty]$ if and only if it is open in $[0, \infty)$.

The first thing we want to prove is that this space is compact. An important fact to note towards this is that every neighborhood U of ∞ contains a set $(a, \infty]$, where $a \in [0, \infty)$. We also want to note that every set that is compact in $[0, \infty)$ is compact in $[0, \infty]$. This is a consequence of the fact that an open set in $[0, \infty]$ that contains ∞ will also be an open set in $[0, \infty)$ after removing ∞. Now, suppose that $\mathcal{U}$ is an open cover of $[0, \infty]$. This $\mathcal{U}$ contains a neighborhood U like we just described, then by our earlier remark $[0, \infty] - U$ will be closed and bounded, and thus compact in $[0, \infty)$. Then $\mathcal{U}$ will have a finite subcover $\mathcal{U}'$ that covers $[0, \infty] - U$. This means that $\mathcal{U}' \cup \{U\}$ is a finite cover of $[0, \infty]$.

We shall actually be looking at the function space $[0, \infty]^X$ for some set X. Any function space Y^X can be identified with the product space $\prod_{x \in X} Y$. This way, if Y is a topological space, Y^X is naturally endowed with the product topology. An important consequence for this topological space Y^X is that Tychonoff's theorem says that it is compact iff Y is compact. The interesting thing for us, of course, is that this implies that $[0, \infty]^X$ is compact.

There is one more result in topology we want to prove before we the next theorem. It is the following alternative definition of compactness.

Lemma 4.1. *A topological space X is compact if and only if every collection of closed subsets* *of X satisfying the* **finite intersection property** *(which says that the intersection of every finite subcollection is nonempty) has nonempty intersection itself. Note, when talking about the intersection of a collection $\mathcal{M}$, we mean $\bigcap_{M \in \mathcal{M}} M$.*

Proof. Let's first suppose that X is compact and that $\mathcal{M} = \{M_i \mid i \in I\}$ is a collection of closed subsets of X that has the finite intersection property. We will argue by contradiction. So, suppose that the sets in $\mathcal{M}$ do have empty intersection. Then the collection $\mathcal{M}' = \{M_i^c \mid i \in I\}$ (where the upper index c indicates the complement of a set) is an open cover of X. Compactness says that there exists a finite subcover of $\mathcal{M}'$, let this subcover be indexed by the finite index set $J \subseteq I$. Then the finite subcollection $\{M_j \mid j \in J\} \subseteq \mathcal{M}$ has empty intersection, this contradicts $\mathcal{M}$ having the finite intersection property.

For the other direction, first suppose that every collection of nonempty subsets of X that satisfies the finite intersection property has nonempty intersection. We shall again argue by contradiction. So, suppose that X is not compact and thus that there exists an open cover $\mathcal{M} = \{M_i \mid i \in I\}$ that doesn't admit a finite subcover. Using the same notation as before, let $\mathcal{M}' = \{M_i^c \mid i \in I\}$. All sets in $\mathcal{M}'$ are closed and $\mathcal{M}'$ also admits the finite intersection property. The latter because, if some finite collection $\{M_j^c \mid j \in J\} \subseteq \mathcal{M}'$ were to have empty intersection, the collection $\{M_j^c \mid j \in J\} \subseteq \mathcal{M}$ would be a finite subcover, which would imply that X is compact. Now, our supposition combined with these two properties of $\mathcal{M}'$ imply that it has nonempty intersection. This however, means that $\mathcal{M}$ does not cover X, this is a contradiction.

□

Theorem 4.2. *Let $(\mathcal{T}, +)$ be a commutative semigroup with identity 0 and let $\varepsilon \in \mathcal{T}$ be fixed. Then $(n+1)\varepsilon \not\leq n\varepsilon$ for all $n \in \mathbb{N}$ if and only if there exists a function $\mu : \mathcal{T} \to [0, \infty]$ that normalizes ϵ (i.e. $\mu(\varepsilon) = 1$) and is finitely additive .*

Proof. Let's start with assuming that a measure μ as described exists and that there is an $n \in \mathbb{N}$ such that $(n+1)\varepsilon \leq n\varepsilon$. This would imply that $n + 1 = \mu((n+1)\varepsilon) \leq \mu(n\varepsilon) = n$ which is an obvious contradiction.

For the other direction assume that $(n+1)\varepsilon \not\leq n\varepsilon$ for all $n \in \mathbb{N}$. First note that if we have a measure like described on the set of bounded (with respect to ε) elements in $\mathcal{T}$, we can extend this measure to all of $\mathcal{T}$ by letting the measure be infinite on the unbounded elements. We shall therefore assume from now on that all elements in $\mathcal{T}$ are bounded. We shall now state and prove a claim that will form the core of our proof.

Claim 1. *Suppose that $\mathcal{T}'$ is a finite subset of $\mathcal{T}$ that contains ε. Then there exists a function $\mu : \mathcal{T}' \to [0, \infty]$ that has the following two properties:*

(a) *μ normalizes ε*

(b) *if $\alpha_1, ..., \alpha_n$ and $\beta_1, ..., \beta_m$ are sequences in $\mathcal{T}'$ such that $\alpha_1 + ... + \alpha_n \leq \beta_1 + ... + \beta_m$, then $\sum \mu(\alpha_i) \leq \sum \mu(\beta_j)$.*

Note that in (b) we only wrote a summation sign when adding numbers, we shall continue doing this in proof for reasons of clarity.

Proof. We shall use induction on the size of $\mathcal{T}'$. Firstly, if $|\mathcal{T}'| = 1$ it will only contain ε. Property (a) tells us that $\mu(\varepsilon) = 1$, in this case this defines the entire function. Because of the premise of the main theorem we have that $k\varepsilon \leq n\varepsilon$ if and only if $k \leq n$ (for if $k > n$ and $k\varepsilon \leq n\varepsilon$ we would have $(n+1)\varepsilon \leq k\varepsilon \leq n\varepsilon$). With this it can easily be seen that (b) holds for μ.

Now suppose that $|\mathcal{T}'| > 1$ and that our claim is true for similar sets of smaller size. Let τ be any element in $\mathcal{T}' - \{\varepsilon\}$. Then, by our induction hypothesis there exists a function $\nu : \mathcal{T}' - \{\tau\} \to [0, \infty]$ that satisfies both (a) and (b). The first thing to note about ν is that it only takes on finite values, we can show this as follows. Let α be an element in $\mathcal{T}' - \{\tau\}$ then our earlier assumption regarding boundedness of all elements in $\mathcal{T}$ tells us that there exists a $k \in \mathbb{N}$ such that $\alpha \leq k\varepsilon$, then directly applying (b) and then (a) tells us that $\nu(\alpha) \leq k\nu(\varepsilon) = k$. We now want to extend this ν to all of $\mathcal{T}'$ without destroying one of its properties (a) or (b), we shall do this by assigning an appropriate value to τ.

So, let μ agree with ν on $\mathcal{T}' - \{\tau\}$ and define $\mu(\tau) = \inf\{(\sum \nu(\gamma_c) - \sum \nu(\eta_d))/r\}$, where the inf is taken over all sequences $\gamma_1, ..., \gamma_p$, $\eta_1, ..., \eta_q$ in $\mathcal{T}' - \{\tau\}$ and $r \in \mathbb{N} - 0$ for which $\eta_1 + ... + \eta_q + r\tau \leq \gamma_1 + ... + \gamma_p$. We will have to note a few facts to show that this choice of $\mu(\tau)$ is valid. Firstly, because τ is bounded, letting $\eta_1, ..., \eta_q$ be empty, $r = 1$, $\gamma_c = \varepsilon$ for all c and p large enough, shows that the inf is taken over a nonempty set and will therefore exist. It also directly shows that $\mu(\tau)$ will not be infinite. Also, for any $\gamma_1, ..., \gamma_p$, $\eta_1, ..., \eta_q$ and r, if $\eta_1 + ... + \eta_q + r\tau \leq \gamma_1 + ... + \gamma_p$ then $\eta_1 + ... + \eta_q \leq \gamma_1 + ... + \gamma_p$. Applying (b) on this last inequality tells us that $\sum \nu(\gamma_c) - \sum \nu(\eta_d) \geq 0$. So, the inf described above will also be greater or equal than zero and thus, μ still only takes on values in $[0, \infty]$.

That μ still satisfies (a) follows directly from its definition. That the same is true for (b) will require a bit more work. To show that (b) still holds we will have to show that for all sequences $\alpha_1, ..., \alpha_n$, $\beta_1, ..., \beta_m$ in $\mathcal{T}' - \{\tau\}$ and $s, t \in \mathbb{N}$, the inequality $\alpha_1 + ... + \alpha_n + s\tau \leq \beta_1 + ... + \beta_m + t\tau$ implies that $\sum \nu(\alpha_i) + s\mu(\tau) \leq \sum \nu(\beta_j) + t\mu(\tau)$. We will prove this for separate values of s and t.

First suppose that $s = t = 0$, in this case the desired inequality follows directly from the fact that ν satisfies (b). Now suppose that $s = 0 < t$. We now have to show that $(\sum \nu(\alpha_i) - \sum \nu(\beta_j))/t \leq \mu(\tau)$. We know that $\mu(\tau)$ is defined as an inf, this means that the inequality still holds if we can show that it is true for an arbitrary upper bound of $\mu(\tau)$. We want to use this fact, so let $(\sum \nu(\gamma_c) - \sum \nu(\eta_d))/r$ be such an upper bound (with this we mean that $\eta_1 + ... + \eta_q + r\tau \leq \gamma_1 + ... + \gamma_p$). We then note the following:
We know that

$$\alpha_1 + ... + \alpha_n \leq \beta_1 + ... + \beta_m + t\tau,$$

multiplying by r and adding $t\eta_1 + ... + t\eta_q$ yields

$$r\alpha_1 + ... + r\alpha_n + t\eta_1 + ... + t\eta_q \leq r\beta_1 + ... + r\beta_m + rt\tau + t\eta_1 + ... + t\eta_q,$$

substituting the inequality regarding the upper bound of $\mu(\tau)$ then gives

$$r\alpha_1 + ... + r\alpha_n + t\eta_1 + ... + t\eta_q \leq r\beta_1 + ... + r\beta_m + t\gamma_1 + ... + t\gamma_p,$$

applying (b) with respect to ν gives

$$r\sum\nu(\alpha_i)+t\sum\nu(\eta_d)\leq r\sum\nu(\beta_j)+t\sum\nu(\gamma_c)$$

because we are now dealing with real numbers and no longer working in a semigroup we are allowed to subtract and divide, this gives

$$\left(\sum\nu(\alpha_i)-\sum\nu(\beta_j)\right)/t\leq\left(\sum\nu(\gamma_c)-\sum\nu(\eta_d)\right)/r$$

which is what we intended to show.

Lastly, suppose that $s>0$, in this case we know that $\alpha_1+...+\alpha_n+s\tau\leq\beta_1+...+\beta_m+t\tau$ and we have to show that $\sum\nu(\alpha_i)+s\mu(\tau)\leq\sum\nu(\beta_j)+t\mu(\tau)$. We are going to use the same reasoning as before, we will show that the inequality holds when the $\mu(\tau)$ on the right hand side is substituted with an arbitrary upper bound, this will imply that the general inequality also holds. So like before, let $(\sum\nu(\gamma_c)-\sum\nu(\eta_d))/r$ be an upper bound of $\mu(\tau)$ (this means that $\eta_1+...+\eta_q+r\tau\leq\gamma_1+...+\gamma_p$).
Like we said before, we know that

$$\alpha_1+...+\alpha_n+s\tau\leq\beta_1+...+\beta_m+t\tau,$$

multiplying by r and adding $t\eta_1+...+t\eta_q$ gives

$$r\alpha_1+...+r\alpha_n+t\eta_1+...+t\eta_q+rs\tau\leq r\beta_1+...+r\beta_m+rt\tau+t\eta_1+...+t\eta_q,$$

substituting the inequality regarding the upper bound of $\mu(\tau)$ gives

$$r\alpha_1+...+r\alpha_n+t\eta_1+...+t\eta_q+rs\tau\leq r\beta_1+...+r\beta_m+t\gamma_1+...+t\gamma_p,$$

looking back at the definition of $\mu(\tau)$ we note that this last inequality implies that

$$\mu(\tau)\leq(r\sum\nu(\beta_j)+t\sum\nu(\gamma_c)-r\sum\nu(\alpha_i)-t\sum\nu(\eta_d))/rs,$$

multiplying by s and adding $\sum\nu(\alpha_i)$ gives

$$\sum\nu(\alpha_i)+s\mu(\tau)\leq\sum\nu(\beta_j)+t(\sum\nu(\gamma_c)-\sum\nu(\eta_d))/r,$$

which is the inequality we wanted to show. This ends the proof of the claim.

Let $\mathcal{T}'$ be as in the claim, we define $\mathcal{M}(\mathcal{T}')\subseteq[0,\infty]^{\mathcal{T}}$ to be the set that contains all functions $\nu:\mathcal{T}\to[0,\infty]$ that normalize ε and are finitely additive on $\mathcal{T}'$. The first thing we want to note about $\mathcal{M}(\mathcal{T}')$ is that our claim implies that it is nonempty. We can construct an appropriate function as follows. Let $\mu:\mathcal{T}'\to[0,\infty]$ be a function like described in the claim, this μ obviously normalizes ε. That μ is also finitely additive can be seen by applying property (b) of the claim twice, note here that we are allowed to apply (b) on α,β and in particular $\alpha+\beta$ because they are all assumed to be inside $\mathcal{T}'$ when considering finite additivity. Now, any function $\nu:\mathcal{T}\to[0,\infty]$ that agrees with this μ on $\mathcal{T}'$ will be an element in $\mathcal{M}(\mathcal{T}')$.

Next we want to show that $\mathcal{M}(\mathcal{T}')$ is closed and that the collection of all $\mathcal{M}(\mathcal{T}')$'s has the finite intersection property. We will start with closedness, one way to prove this is by showing that the complement of $\mathcal{M}(\mathcal{T}')$ is open. We will do this by showing that every element outside of $\mathcal{M}(\mathcal{T}')$ has a neighborhood that is disjoint from $\mathcal{M}(\mathcal{T}')$. So, consider a function $\nu : \mathcal{T} \to [0, \infty]$ that isn't in $\mathcal{M}(\mathcal{T}')$, this means that ν doesn't normalize ε or isn't finitely additive. Let's first suppose that $\nu(\varepsilon) \neq 1$, then, by the product topology, the set $\{\nu' : \mathcal{T} \to [0, \infty] \mid \nu'(\varepsilon) \in (\nu(\varepsilon) - \delta, \nu(\varepsilon) + \delta)\}$ is an open neighborhood of ν in $[0, \infty]^{\mathcal{T}}$ and when δ is chosen small enough, all functions in this set will not normalize ε, which means that this set will be disjoint from $\mathcal{M}(\mathcal{T}')$. We can apply the same idea in the situation where ν isn't finitely additive. If α, β are elements such that $\nu(\alpha) + \nu(\beta) \neq \nu(\alpha + \beta)$, then, again by the product topology, the set $\{\nu' : \mathcal{T} \to [0, \infty] \mid \nu'(\alpha) \in (\nu(\alpha) - \delta, \nu(\alpha) + \delta)$ and $\nu'(\beta) \in (\nu(\beta) - \delta, \nu(\beta) + \delta)\}$ is an open neighborhood of ν and when δ is chosen small enough, all functions in this set will fail to be finitely additive on α, β.

For the finite intersection property, suppose that $\mathcal{T}_1', ..., \mathcal{T}_k'$ are all like $\mathcal{T}'$ as in the claim. We then want to show that $\mathcal{M}(\mathcal{T}_1') \cap ... \cap \mathcal{M}(\mathcal{T}_k')$ is nonempty. For this, first note that a function that is finitely additive on some set will also be finitely additive on any subset of this set. This means that we can write $\mathcal{M}(\mathcal{T}_1' \cup ... \cup \mathcal{T}_k') \subseteq \mathcal{M}(\mathcal{T}_1') \cap ... \cap \mathcal{M}(\mathcal{T}_k')$. Then note that $\mathcal{T}_1' \cup ... \cup \mathcal{T}_k'$ still has the same properties as $\mathcal{T}'$ from the claim, we argued earlier that this implies that $\mathcal{M}(\mathcal{T}_1' \cup ... \cup \mathcal{T}_k')$ is nonempty.

We have now shown that $\{\mathcal{M}(\mathcal{T}') \mid \mathcal{T}'$ is finite and contains ε $\}$ contains only closed sets and satisfies the finite intersection property. Because this collection is defined on the compact space $[0, \infty]^{\mathcal{T}}$, lemma 4.1 tells us that the intersection of this collection is nonempty. Now any function μ in this intersection normalizes ε and is finitely additive on any $\mathcal{T}'$. That it is also finitely additive on $\mathcal{T}$ can be seen by by taking $\mathcal{T}' = \{\varepsilon, \alpha, \beta, \alpha + \beta\}$ for any $\alpha, \beta \in \mathcal{T}$.

□

This theorem finally enables us to prove the reverse direction of theorem 1.9 .

Theorem 4.3. *(Tarski) Let G be a group acting on a set X and suppose $E \subseteq X$. Then, if E is not paradoxical, there exists a finitely additive measure $\mu : \mathcal{P}(X) \to [0, \infty]$ that normalizes E and is also G-invariant.*

Proof. Let $\mathcal{S}$ be the type semigroup induced by the action of G on X and denote $\varepsilon = [E \times \{0\}]$. Then, because E is not paradoxical remark 3.4 tells us that $2\varepsilon \neq \varepsilon$, applying corollary 3.12 then says that $(n+1)\varepsilon \neq n\varepsilon$ for all $n \in \mathbb{N}$. This means that we can apply theorem 4.2, which gives us a finitely additive function $\nu : \mathcal{S} \to [0, \infty]$ that normalizes ε. We then use this ν to define the desired measure on X. For $A \in \mathcal{P}(X)$ let $\mu(A) = \nu([A \times \{0\}])$, the construction of this μ does not destroy any of the previously mentioned properties of ν and is also G-invariant since $A \times \{0\} \sim gA \times \{0\}$. □

The reverse direction of theorem 1.10 is now an easy consequence of our last theorem.

Theorem 4.4. *A group that is not paradoxical, is amenable.*

5 Left invariant means on locally compact groups

In this chapter we will take a look at another equivalent definition of amenability that we are able to state after endowing groups with the discrete topology. We will then generalize this new definition for groups with a less specific topology endowed upon them. After we have done this we want to prove an interesting hereditary property regarding this generalized notion of amenability, the majority of this chapter will be spent proving this result. This result also allows us to prove a property regarding the heredity of amenability with respect to the the subgroup relation. Note, however, that this is just an easy consequence of the main results of this chapter, there are far easier ways to prove this hereditary property (which we will not discuss here). Most results in this chapter are a bit less elementary than the previous chapters, because of this we will assume some basic knowledge in measure theory and functional analysis, and we will also omit some (parts of some) proofs because their methods are not very relevant and/or overly tedious. The first thing we need to do is introduce the notion of topological groups.

Definition 5.1. *A* **topological group** *is a group G which is also a topological space such that the product and inverse functions on G*

$$G \times G \to G, \quad (x, y) \to xy \quad \textit{and} \quad G \to G, \quad x \to x^{-1}$$

are continuous.

Any group G is a topological group when endowed with the discrete topology. We call such G discrete groups. A particular kind of topological groups we will be looking at are the **locally compact groups**, these are topological groups that are also locally compact, Hausdorff spaces (note that any discrete group is locally compact). An important result regarding locally compact groups is the following theorem, which we will state without proof and accompanied by a remark about some properties regarding the theorem.

Theorem 5.2. *Let G be a locally compact group, then G is endowed with a non-zero, locally finite (i.e. every point in g has an open neighborhood of finite measure) , regular (i.e. both inner and outer regular), non-negative Borel measure that is left-invariant on all Borel subsets of G, we denote this measure with m_G.*

Remark 5.3. *Let G be as above, then m_G has the following properties:*

(a) For every compact $K \subseteq G$ we have that $m_G(K) < \infty$

(b) For every nonempty, open $U \subseteq G$ we have that $m_G(U) > 0$

Proof. For (a), suppose that $K \subseteq G$ is compact and that $m_G(K) = \infty$, by local finiteness there exists a neighborhood U of the identity element such that U has finite measure. Then the collection $\{kU \mid k \in K\}$ is an open cover of K, by compactness there exists a finite subcover. By finite additivity and left invariance of m_G this implies that there is an $n \in \mathbb{N}$ such that $n\, m_G(U) \geq m_G(K) = \infty$, which implies that $m_G(U) = \infty$, which is a contradiction.

For (b), suppose without loss of generality that U is an open neighborhood of 0 such that $m_G(U) = 0$ and let K be an arbitrary compact subset of G. Then the collection $\{kU \mid k \in$

$K\}$ is an open cover of K, by compactness there exists a finite subcover. By finite additivity and left invariance of m_G this implies that there is an $n \in \mathbb{N}$ such that $m_G(K) \leq nm_G(U) = 0$, which implies that $m_G(K) = 0$. Then, using inner regularity we get that $m_G(G) = \sup\{m_G(K) \mid K \text{ is compact}\} = 0$, but this is in contradiction with the assumption that m_G is non-zero. $\square$

When G is locally compact and we talk about $L^\infty(G)$, it is with respect to the above measure. In the discrete case, we suppose that m_G is the counting measure and we then get that $L^\infty(G) = \ell^\infty(G)$. We are now ready to state the definition of mean, after which we are able to give the equivalent definition of amenability. Along with the definition of a mean we shall state a theorem (without proof) regarding means that we will use later on in this chapter.

Definition 5.4. *Let G be a locally compact group and E a subspace of $L^\infty(G)$ that contains all constant functions and is closed under complex conjugation. We call a linear functional $M \in E^*$ (where E^* is the dual space of E) a* **mean** *if $M(1) = ||1|| = 1$ (note that the symbol 1 is used ambiguously). Moreover, we call E* **left invariant** *if for every $\phi \in E$ and $g \in G$ we have that ${}_g\phi(\cdot) = \phi(g^{-1}\cdot) \in E$ too. For such E we call a mean M* **left invariant** *if for every $\phi \in E$ and $g \in G$ we have that $M({}_g\phi) = M(\phi)$.*

Theorem 5.5. *Let G be a locally compact group. There exists a left invariant mean on $L^\infty(G)$ if and only if there exists one on $\mathcal{C}_b(G)$, the space of continuous, bounded (complex) functions. Note that this makes sense because, by continuity of complex conjugation, and product and inverse functions (on G) we have that $\mathcal{C}_b(G)$ is both closed under complex conjugation and left-invariant.*

Theorem 5.6. *A discrete group G is amenable if and only if there exists a left invariant mean on $L^\infty(G)$.*

Proof. Lets first suppose that G is amenable, this gives us a finitely additive measure μ on $\mathcal{P}(G)$ (the topology of the discrete group) that normalizes G and is left invariant. Without delving into any details, we can apply the usual construction of integrals on this measure even though it is only finitely additive (in contrast to being σ-additive). The functional given by

$$L^\infty(G) \to \mathbb{C}, \quad \phi \to \int_G \phi \, d\mu$$

then defines a left invariant mean.

For the other direction suppose that we are given a left invariant mean M on $L^\infty(G)$. Then define the function $\mu : \mathcal{P}(G) \to [0, \infty]$ as follows, for any $A \subseteq G$ let $\mu(A) = M(\chi_A)$, where χ denotes the indicator function on G. This μ inherits finite additivity from χ (when regarded as a measure on $\mathcal{P}(X)$ instead of a function on G) and left invariance from M. So we are done. $\square$

Our next aim is to show that, for a locally compact group the existence of a left invariant mean like described above is hereditary with respect to closed subgroups. We shall prove that this is true by using Bruhat functions, the definition of which we shall give later on. Before we can prove this we will need to prove about a dozen related statements and introduce some definitions. We will work towards the desired statement as linearly as possible, and accompany the result with some of its interesting consequences.

Definition 5.7. *Let X be a topological space. We call an element $x \in X$ a* **cluster point** *of an infinite sequence $(x_n)_{n\in\mathbb{N}}$ if for every open neighborhood U of x there are infinitely many $n \in \mathbb{N}$ such that $x_n \in U$.*

Proposition 5.8. *Let X be a compact topological space, then every infinite sequence $(x_n)_{n\in\mathbb{N}}$ in X has a cluster point in X.*

Proof. Suppose that $(x_n)_{n\in\mathbb{N}}$ doesn't have a cluster point, then every $x \in X$ has an open neighborhood U_x for which there exist only finitely many $n \in \mathbb{N}$ such that $x_n \in U_x$. We use all such neighborhoods to construct an open cover of X, let $\mathcal{U} = \{U_x \mid x \in X\}$. Then, by compactness there exists a finite subset $X' \subseteq X$ such that $\mathcal{U}' = \{U_x \mid x \in X'\}$ still covers X. However, by construction, there are only a finite amount of elements of the sequence (x_n) inside every $U_x \in \mathcal{U}'$, therefore, $\mathcal{U}'$ does not cover the entire infinite sequence (x_n) and can't be a cover of X, this is a contradiction. $\square$

Proposition 5.9. *Let G be a topological group, if $U \subseteq G$ is symmetric (i.e. $U = U^{-1}$), then the topological closure of U and the interior of U are both symmetric too.*

Proof. The first thing to note is that the inverse function is actually a homeomorphism on G, it is obviously a bijection, and its inverse is the same function and is therefore also continuous. That the closure of U is symmetric can be seen as follows. If A is a closed set that contains U, the set $A \cap A^{-1}$ is symmetric, closed and still contains U (by U's symmetry). Therefore, the closure of U, which is defined as the intersection of all closed sets that contain U is also symmetric. That the interior is symmetric follows from more or less the same reasoning. If $x \in \mathring{U}$ there exists an open $U_x \subseteq U$ that contains x. The set U_x^{-1} is an neighborhood of x^{-1}, it is open by our remark on the inverse function, and it is a subset of U because U is symmetric. Therefore $x^{-1} \in \mathring{U}$ which implies that $\mathring{U}$ is symmetric. $\square$

Proposition 5.10. *Let G be a topological group, if U is an open neighborhood of the identity element, there exists a symmetric open neighborhood V of the identity such that $VV \subseteq U$.*

Proof. We will use the fact that the multiplication function is continuous. This, together with the product topology tells us that there exists an open set $V_1 \times V_2 \subseteq G \times G$ such that $V_1V_2 = U$. Also, by the product topology both V_1 and V_2 are open in G, and moreover, we can assume that they are both neighborhoods of the identity element $e \in G$ since the preimage of any set containing e under the multiplication function will contain $\{e\} \times \{e\}$. Now, define $V_3 = V_1 \cap V_2$, this open neighborhood of e has the property that $V_3V_3 \subseteq V_1V_2 = U$. Finally, let $V = V_3 \cap V_3^{-1}$, since e is its own inverse and the inverse function is a homeomorphism this is open neighborhood of e, this set is also symmetric and still has the property that $VV \subseteq U$. So, this V has all the desired properties. $\square$

Lemma 5.11. *Let G be a locally compact group. If H is a closed subgroup of G and U an open, symmetric and* **relatively compact** *(i.e. the topological closure of U is compact) neighborhood of the identity element e, there exists a subset A of G that has the following two properties:*

(a) For every $g \in G$ there exists a in A such that the intersection $gH \cap Ua$ is nonempty

(b) For every compact $K \subseteq G$ there are only finitely many $a \in A$ such that the intersection $KH \cap \overline{U}a$ is nonempty

Proof. We will first apply Zorn's lemma to find an $A \subseteq G$ such that $a \notin Ua'H$ for all $a \neq a'$ in A that is maximal with regard to set inclusion. After this we shall prove that this A satisfies (a) and (b).

Consider the poset of subsets B of G that have the property that $b \notin Ub'H$ for every $b \neq b'$ in B, order this poset by set inclusion. Note that the empty set is an element of this poset because it satisfies all for all statements, this means that the empty chain has an upper bound. Now, suppose that $\{B_i \mid i \in I\}$ is a nonempty chain in the poset and let $B = \bigcup_{i \in I} B_i$. This B is obviously an upper bound if it is an element in the poset. To show this is the case, let $b \neq b'$ be elements in B, because of the definition of B there exist $i, i' \in I$ such that $b \in B_i$ and $b' \in B_{i'}$. Moreover, because all the B_i's form a chain, we can say without loss of generality that $B_i \subseteq B_{i'}$. Therefore b and b' are both elements in $B_{i'}$ (which is an element in the poset) and have the property that $b \notin Ub'H$. Zorn's lemma now implies the existence of a maximal element in the poset, denote this element with A.

We now want to prove that A satisfies property (a), we will do this using contradiction. Suppose that g is an element in G such that $gH \cap Ua = \emptyset$ for all $a \in A$. This is equivalent with saying that for all $a \in A$ and $h \in H$ we have $gh \notin Ua$, which is equivalent with saying that for all $a \in A$ and $h \in H$ we have $g \notin Uah^{-1}$, which, since H is a group (and thus symmetric), is equivalent with saying that $g \cap UaH = g \cap UaH^{-1} = \emptyset$ for all $a \in A$. Because of U's symmetry we can apply the same reasoning to find that $a \notin UgH$ for all $a \in A$. Now, A's maximality implies that $g \in A$, for if this were not the case, the set $A \cup \{g\}$ would contradict A's maximality. However, since U and H both contain the identity element, we have that $gH \cap Ug \neq \emptyset$, which contradicts our assumption.

For the second property we will also argue by contradiction. Suppose that K is a compact subset of G with the property that there are infinitely many $a \in A$ such that $KH \cap \overline{U}a$ is nonempty. Then, using the fact that H and $\overline{U}$ are symmetric and the same reasoning as before, we know that there infinitely many $a \in A$ such that $\overline{U}K \cap aH \neq \emptyset$. This tells us that there exists a sequence $(a_i)_{i \in \mathbb{N}}$ in A, where all a_i are distinct and a sequence $(h_i)_{i \in \mathbb{N}}$ in H such that $(a_i h_i)_{i \in \mathbb{N}}$ is a sequence in $\overline{U}K$. Combining the facts that the product function on G is continuous and that both $\overline{U}$ and K are compact implies that $\overline{U} \times K$ is compact in the product topology, we get that $\overline{U}K$ is compact too. Then, proposition 5.8 tells us that $(a_i h_i)$ has a cluster point g in $\overline{U}K$. Now, we can use proposition 5.10 to find a symmetric open neighborhood V of e such that $VV \subseteq U$. Then, Vg is an open neighborhood of the cluster point g, therefore there exist $n \neq m$ in $\mathbb{N}$ such that $a_n h_n$ and $a_m h_m$ are both elements in Vg. Then, $(a_n h_n)(a_m h_m)^{-1} \in Vg(Vg)^{-1} = Vgg^{-1}V^{-1} = VV \subseteq U$, or $a_n \in U a_m h_m h_n^{-1}$, this then implies that $a_n \in U a_m H$, which is impossible because of the definition of A and the fact that all a_i are distinct, □

Proposition 5.12. *Let X be a locally compact, Hausdorff space and let $x \in X$. For every neighborhood U of x there exists a relatively compact open neighborhood V of x such that $\overline{V} \subseteq U$.*

Proof. Using X's local compactness we get a compact neighborhood K of x. Let W be the interior of $K \cap U$, this is an open neighborhood of x. Then $K - W$ is compact since it is closed and a subset of K. For every $k \in K - W$ apply Hausdorffness to obtain an open neighborhood A_k of k that doesn't contain x. The collection of all such neighborhoods forms an open cover of $K - W$, using $K - W$'s compactness gives us a finite subcover $\{A_{k_1}, ..., A_{k_n}\}$. Then, let

$V' = K - (A_{k_1} \cup ... \cup A_{k_n})$, this is a closed neighborhood of x and a subset of both U and K. Finally, let $V = \overset{\circ}{V'}$, then V is an open neighborhood of x, and since its closure is a subset of V' it will also be a subset of U like we want it to and be a subset of K which makes V relatively compact. ☐

Proposition 5.13. *Any locally compact, Hausdorff space X is regular.*

Proof. Suppose that $A \subseteq X$ is closed and that we are given an $x \notin A$. By proposition 5.12 there exists an open neighborhood U of x such that $U \subseteq \overline{U} \subseteq A^c$. Then $\overline{U}^c$ is a(n) (open) neighborhood of A that doesn't intersect with the neighborhood U of x, this proves regularity. ☐

Proposition 5.14. *Any topological space X that is σ-***compact** *(i.e. a countable union of compact sets) is also a* **Lindelöf** *space (i.e. every open cover has a countable subcover).*

Proof. Suppose that $X = \bigcup_{i\in\mathbb{N}} A_i$ where A_i is compact for all $i \in \mathbb{N}$ and let $\mathcal{U}$ be an open cover of X. Because $\mathcal{U}$ also covers every A_i, there is, for every $i \in \mathbb{N}$, a finite subcover $\mathcal{U}_i \subseteq \mathcal{U}$ of A_i. Then the collection $\bigcup_{i\in\mathbb{N}} \mathcal{U}_i$ countable union of finite collections and is countable itself, it covers every A_i and therefore also covers X, this proves that X is a Lindelöf space. ☐

Proposition 5.15. *If a topological space X is partitioned by subspaces that are both normal and closed, X itself is normal too.*

Proof. Denote the collection that partitions X by $\{A_i \mid i \in I\}$ where I is an indexing set. To prove normality of X, suppose that E and F are closed, disjoint subsets of X. For every indexing element $i \in I$ let $E_i = A_i \cap E$ and $F_i = A_i \cap F$, this way the collections $\{E_i \mid i \in I\}$ and $\{F_i \mid i \in I\}$ partition E and F respectively. For fixed i we also have that E_i and F_i are both closed, disjoint subsets of A_i. Applying normality of A_i then gives disjoint neighborhoods $U_i \subseteq A_i$ and $V_i \subseteq A_i$ of E_i and F_i respectively. Then, the unions $\bigcup_{i\in I} U_i$ and $\bigcup_{i\in I} V_i$ are disjoint neighborhoods of E and F respectively, this proves normality of X. ☐

Theorem 5.16. *Any locally compact group G is a normal topological space.*

Proof. We will first prove the following short claim

Claim 1. *For any symmetric neighborhood V of the identity element e we have that $\overline{V} \subseteq VV$*

proof. Let $x \in \overline{V}$, then xV is a neighborhood of x. This tells us that the intersection $V \cap xV$ is nonempty. Suppose that y is an element in this intersection, we can write $y = xz$ for a $z \in V$, but then we can also write $x = yz^{-1} \in VV^{-1} = VV$, since this is what we intended to show this ends the proof of the claim.

Now, let K be compact neighborhood of the identity element and then let V be the interior of $K \cap K^{-1}$, then V is a symmetric (by proposition 5.9), relatively compact, open neighborhood of e (for relative compactness use that K is compact inside a Hausdorff space and is therefore closed, and that this means that $\overline{V}$ is closed inside the compact set K). Now, let $H = \bigcup_{n\geq 1} V^n$, then H is a subgroup of G, open, and also σ-compact (and therefore Lindelöf) since the continuity of the product function on G implies that $\overline{V}^n$ is compact for each $n \geq 1$ and we have that $H = \bigcup_{n\geq 1} V^n = \bigcup_{n\geq 1} \overline{V}^n$ by the claim. Because any locally compact group is regular and any regular Lindelöf space is normal (for this result we refer to a general topology book), we get that H is normal. Note that the continuity of the product function also implies that left and right translation are homeomorphisms, this tells us that all the cosets of H are also open. Since the cosets of H partition G, we have that H is closed. Again, using that all the cosets of H are homeomorphic to H we find a partition of G that consists of subspaces that are both normal and closed, by proposition 5.15 this implies that G itself is normal too. □

Definition 5.17. *Let X be a topological space and $f : X \to \mathbb{C}$ a function, we define the* **support** *of f as the following set:* $\mathrm{supp}(f) = \overline{\{x \in X \mid f(x) \neq 0\}} = \overline{f^{-1}(\{0\}^c)}$. *We denote the set of continuous functions on X with compact support as $\mathcal{C}_{00}(X)$.*

Proposition 5.18. *Let G be a locally compact group, then there exists a non-negative, real function $\phi \in \mathcal{C}_{00}(G)$ that normalizes the identity element e and such that $\phi(g) = \phi(g^{-1})$ for all $g \in G$.*

Proof. Applying proposition 5.12 on G gives us us an open neighborhood $U \subseteq G$ of e such that $\overline{U}$ is compact, applying it again on U gives us a compact (and therefore closed) neighborhood K of e that is a subset of U. Applying Urysohn's lemma on K and the complement of U gives us a continuous function $\psi : G \to [0, 1]$ such that $\phi\big|_K = 1$ and $\phi\big|_{U^c} = 0$. This ψ normalizes e because $e \in K$, and ψ has compact support because $\mathrm{supp}(f)$ is a closed subset of the compact set $\overline{U}$. Then, let $\phi(g) = (\psi(g) + \psi(g^{-1}))/2$, the construction of this function doesn't destroy any of the properties of ψ we mentioned above but does create the extra property that $\phi(g) = \phi(g^{-1})$. Therefore, ϕ's existence proves our statement. □

Lemma 5.19. *Let G be a locally compact group. For every closed subgroup $H < G$ there exists a continuous function $\phi : G \to [0, \infty)$ that has the following two properties:*

(a) For every $g \in G$ the intersection $\{x \in G \mid \phi(x) > 0\} \cap gH$ is nonempty

(b) For every compact $K \subseteq G$, the support $\mathrm{supp}(\phi\big|_{KH})$ is also compact

Proof. Choose a function $\psi \in \mathcal{C}_{00}(G)$ like described in proposition 5.18 and let $U = \{g \in G \mid \psi^{-1}(0, \infty)\}$. We will state some properties of this U. Firstly, since ψ normalizes the identity element e we have that U is a neighborhood of e, also, ψ's continuity tells us that U is open. The fact that $\psi(g) = \psi(g^{-1})$ for all $g \in G$ implies that U is symmetric, and finally, U is relatively compact because $\overline{U} = \mathrm{supp}(\psi)$.

This set of properties imply that we can apply lemma 5.11 , denote the subset of G this lemma gives us with A and define the function $\phi : G \to [0, \infty)$ by $\phi(g) = \sum_{a \in A} \psi(ga^{-1})$. We want to show that this ϕ is a well-defined function and has all the properties we described above. First, let $g \in G$ and use G's local compactness to find a compact neighborhood K of g. Note that since $g \in KH$ (this because $g \in K$ and $e \in H$) property (b) of lemma 5.11 says that there are only finitely many $a \in A$ such that $ga^{-1} \in \overline{U}$, then note that $ga^{-1} \notin \overline{U} = \mathrm{supp}(\psi)$ implies that $\phi(ga^{-1}) = 0$. So, the sum by which $\phi(g)$ is defined is finite for all $g \in G$, this tells us that ϕ is both well-defined and continuous.

For (a), first note that for an $x \in G$ we have that $\phi(x) > 0$ if and only if there exists an $a \in A$ such that $\psi(xa^{-1}) > 0$. Now, let $g \in G$ and use lemma 5.11(a) to find an $a_0 \in A$ and $h \in H$ such that $gha_0^{-1} \in U$. By definition of U we have that $\psi(gha_0^{-1}) > 0$, which implies that $\phi(gh) > 0$, which implies that $\{x \in G \mid \phi(x) > 0\} \cap gH$ is nonempty, which is what we needed to show.

For (b), suppose that $K \subseteq G$ is compact and introduce the notation $V = \{x \in KH \mid \phi(x) > 0\}$. By the equivalence we mentioned when showing (a) we get that $\phi(x) > 0$ if and only if $x \in \bigcup_{a \in A} Ua$, we can then write $V = \bigcup_{a \in A} Ua \cap KH$. By lemma 5.11(b) there is a finite sequence $a_1, ..., a_n$ in A such that $KH \cap \overline{U}a$ is empty if $a \in A - \{a_1, ..., a_n\}$, note that for such a the intersection $KH \cap Ua$ will also be empty. We then get that

$$\operatorname{supp}(\phi|_{KH}) = \overline{V} = \overline{\{x \in KH \mid \phi(x) > 0\}} \subseteq \overline{Us_1 \cup ... \cup Us_n} = \overline{Us_1} \cup ... \cup \overline{Us_n}$$

Then, using the facts that left translation is a homeomorphism and that U is relatively compact we get for each $i \leq n$ that the set $\overline{Us_i}$ is compact and that the union over all of them is compact too. Since $\operatorname{supp}(\phi|_{KH})$ is closed by definition, it inherits compactness by the subset relation. $\square$

Lemma 5.20. *Let G be a locally compact group, $H < G$ a closed subgroup and $\phi : G \to \mathbb{C}$ a continuous function such that $\operatorname{supp}(\phi|_{KH})$ is compact for every compact $K \subseteq G$. Then, for every $\psi \in L^\infty(H)$ the following function is continuous*

$$G \to \mathbb{C}, \quad g \to \int_H \psi(h)\phi(g^{-1}h)\ dm_H(h)$$

Proof. We will first show that this function is actually well-defined. Before we do this we will introduce some notation and make a few remarks. Suppose that $g \in G$, let $A = \operatorname{supp}(\phi|_{g^{-1}H})$ and $A' = g^{-1}H - A$. Since $\{g^{-1}\}$ is compact we get that A is compact by assumption. Then, gA is compact too because left translation is a homeomorphism. Note that ϕ's continuity implies that it is bounded on compact sets. Furthermore, gA's compactness implies that it is closed and therefore Borel measurable. Combining all of this with remark 5.3(a) (which says that m_H is finite on compact sets) means that, for every $g \in G$ we can write the following

$$\begin{aligned}
|\int_H \psi(h)\phi(g^{-1}h)\ dm_H(h)| &\leq \int_H |\psi(h)\phi(g^{-1}h)|\ dm_H(h) \\
&\leq \int_H ||\psi||_\infty |\phi(g^{-1}h)|\ dm_H(h) \\
&= ||\psi||_\infty (\int_{gA} |\phi(g^{-1}h)|\ dm_H(h) + \int_{gA'} |\phi(g^{-1}h)|\ dm_H(h)) \\
&\leq ||\psi||_\infty (\int_{gA} \sup_{x \in A}\{|\phi(x)|\}\ dm_H(h) + \int_{gA'} 0\ dm_H(h)) \\
&\leq ||\psi||_\infty \sup_{x \in A}\{|\phi(x)|\}\ m_H(gA) \\
&< \infty.
\end{aligned}$$

Which establishes well-definedness. Before we are able to prove continuity we will have to show a few other facts, to keep this proof from getting too chaotic we shall state and prove these facts in the form of two claims.

Claim 1. *Every $\phi \in \mathcal{C}_{00}$ is left uniformly continuous (i.e. the function $x \to \phi(x^{-1}...)$ is continuous for every $x \in G$).*

Proof. By continuity of both left translation and inversion, we are done if we can show that for every $\varepsilon > 0$ there exists an open neighborhood U of the identity element e such that $|\phi(xg) - \phi(g)| \leq \varepsilon$, for all $x \in U$ and $g \in G$. So, let $\varepsilon > 0$ and introduce the notation $K = \mathrm{supp}(f)$. Using continuity of both right translation and ϕ we can, for every $k \in K$, find an open neighborhood V_k of e such that $|\phi(xk) - \phi(k)| \leq \varepsilon/2$, for all $x \in V_k$. By proposition 5.10 we can, for all such V_k, find a symmetric, open neighborhood U_k of e such that $U_k U_k \subseteq V_k$. Then the collection $\{U_k k \mid k \in K\}$ covers K and by compactness there exists a finite subset $\{k_1, ..., k_n\} \subseteq K$ such that $\{U_{k_i} k_i \mid i \leq n\}$ still covers K. Let $U = \bigcap_{i \leq n} U_{k_i}$, suppose that $x \in U$ is fixed and let $g \in G$, then, in the case that $g \in K$ there is an $i \leq n$ for which $g \in U_{k_i} k_i$, this means that we can write $xg = x(gk_i^{-1})k_i \in U_{k_i} U_{k_i} k_i \subseteq V_{k_i} k_i$, which implies that

$$|\phi(xg) - \phi(g)| \leq |\phi(xg) - \phi(k_i)| + |\phi(k_i) - \phi(g)| \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

In the case that $g \notin K$ and $xg \in K$ we can pick an $i \leq n$ such that $xg \in U_{k_i} k_i$, which also tells us that $g \in x^{-1} U_{k_i} k_i \subseteq U_{k_i} U_{k_i} k_i \subseteq V_{k_i} k_i$, which again allows us to write

$$|\phi(xg) - \phi(g)| \leq |\phi(xg) - \phi(k_i)| + |\phi(k_i) - \phi(g)| \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

Finally, for $x \notin K$ and $xy \notin K$ we have that $|\phi(xy) - \phi(x)| = |0 - 0| = 0$.

Claim 2. *For ϕ like in the proposition, we have that for every $g \in G$ and $\varepsilon > 0$ there exists a compact neighborhood K of g such that*

$$\int_H |\phi(k_1^{-1}h) - \phi(k_2^{-1}h)| \; dm_H(h) \leq \varepsilon$$

for every $k_1, k_2 \in K$.

Proof. Let $g \in G$ and $\varepsilon > 0$ and let K_0 be compact neighborhood of g. Let $A_0 = \mathrm{supp}(\phi|_{K_0^{-1}H})$, this A_0 is compact by definition of ϕ, and by continuity of the product function $K_0 A_0$ is compact too. Since H is closed we also have that $K_0 A_0 \cap H$ is compact, denote this last intersection with K_0'. By remark 5.3(a,b) we have that $0 < m_H(K_0') < \infty$. Now, by combining claim 1 and proposition 5.12 we can find a compact set $K \subseteq K_0$ such that $||\phi(k_1^{-1}...) - \phi(k_2^{-1}...)||_\infty \leq \varepsilon / m_H(K_0')$ for all $k_1, k_2 \in K$. Let $A = \mathrm{supp}(\phi|_{K^{-1}H})$ and $K' = KA \cap H$, by reasoning as before K' is compact, furthermore note that $K \subseteq K_0$ implies that $A \subseteq A_0$, which in turn implies that $K' \subseteq K_0'$ and thus that $0 < m_H(K') \leq m_H(K_0')$. We are finally able to write

$$\begin{aligned}
&\int_H |\phi(k_1^{-1}h) - \phi(k_2^{-1}h)| \; dm_H(h) \\
&\leq \int_{K'} |\phi(k_1^{-1}h) - \phi(k_2^{-1}h)| \; dm_H(h) + \int_{H-K'} |\phi(k_1^{-1}h) - \phi(k_2^{-1}h)| \; dm_H(h) \\
&\leq \int_{K'} ||\phi(k_1^{-1}...) - \phi(k_2^{-1}...)||_\infty \; dm_H(h) + \int_{H-K'} 0 \; dm_H(h) \\
&= ||\phi(k_1^{-1}...) - \phi(k_2^{-1}...)||_\infty \; m_H(K') \\
&\leq \varepsilon
\end{aligned}$$

for all $k_1, k_2 \in K$. This proves the claim.

With these claims, proving continuity is fairly straightforward. First of all note that continuity is trivial if ψ is the zero function, so assume that it is not. Let $g \in G$ and $\varepsilon > 0$ by claim 2 we can choose a (compact) neighborhood K of g such that for every $k_1, k_2 \in K$ we have

$$\int_H |\phi(k_1^{-1}h) - \phi(k_2^{-1}h)| \; dm_H(h) \leq \varepsilon/||\psi||_\infty.$$

Then note that for every $k_1, k_2 \in K$ we have that

$$\begin{aligned}
&|\int_H \psi(h)(\phi(k_1^{-1}h) - \phi(k_2^{-1}h)) \; dm_H(h)| \\
&\leq \int_H |\psi(h)| \; |\phi(k_1^{-1}h) - \phi(k_2^{-1}h)| \; dm_H(h) \\
&\leq \int_H ||\psi||_\infty |\phi(k_1^{-1}h) - \phi(k_2^{-1}h)| \; dm_H(h) \\
&= ||\psi||_\infty \int_H |\phi(k_1^{-1}h) - \phi(k_2^{-1}h)| \; dm_H(h) \\
&\leq \varepsilon.
\end{aligned}$$

Which shows continuity. □

Lemma 5.21. *Let G be a locally compact group and $H < G$ a closed subgroup, then there exists a function $\beta : G \to [0, \infty)$ with the following properties*

(a) For every compact $K \subseteq G$, the support $\mathrm{supp}(\beta\big|_{KH})$ is also compact

(b) For every $g \in G$, we have

$$\int_H \beta(gh) \; dm_H(h) = 1$$

We call a function like β a **Bruhat function** *for H.*

Proof. We start with applying lemma 5.19 denote the function it gives us with ϕ. We then use this ϕ to define the following function

$$\alpha : G \to \mathbb{C}, \quad \alpha(g) = \int_H \phi(gh) \; dm_H(h).$$

Because of lemma 5.19(b) we can apply lemma 5.20 with $\psi = 1$, combining this with the fact that the inversion function is continuous we get that α is well-defined and continuous. By combining lemma 5.19(a) with remark 5.3(b) we get that $\alpha(g) > 0$ for every $g \in G$. This implies that the following function is well defined and continuous

$$\beta : G \to \mathbb{C}, \quad \beta(g) = \frac{\phi(g)}{\alpha(g)},$$

it actually also implies that $\text{supp}(\beta|_{KH}) = \text{supp}(\phi|_{KH})$ which means that β satisfies (a). For (b) note that m_H's left invariance implies that α is constant on the left cosets of H (i.e. $\alpha(gh) = \alpha(g)$ for all $g \in G$). Then, for every $g \in G$ we get

$$\int_H \beta(gh)\ dm_H(h) = \int_H \frac{\phi(gh)}{\alpha(gh)}\ dm_H(h) = \frac{1}{\alpha(g)} \int_H \phi(gh)\ dm_H(h) = \frac{\alpha(g)}{\alpha(g)} = 1.$$

□

Theorem 5.22. *Let G be a locally compact group and $H < G$ a closed subgroup. If there exists a left invariant mean on $L^\infty(G)$, then there exists one on $L^\infty(H)$ too.*

Proof. Lemma 5.21 gives us a Bruhat function $\beta : G \to \mathbb{C}$ for H, we use this β to define the following function

$$T : \mathcal{C}_b(H) \to \ell^\infty(G), \quad (T\psi)(g) = \int_H \psi(h)\beta(g^{-1}h)\ dm_H(h) \quad (g \in G)$$

where T is working on ψ. First note that this T is well-defined since lemma 5.21(b) tells us that

$$\begin{aligned}
|(T\psi)|_\infty &= \sup_{g\in G}\{|\int_H \psi(h)\beta(g^{-1}h)\ dm_H(h)|\} \\
&\leq \sup_{g\in G}\{\int_H |\psi(h)\beta(g^{-1}h)|\ dm_H(h)\} \\
&\leq \sup_{g\in G}\{\int_H ||\psi||_\infty|\beta(g^{-1}h)|\ dm_H(h)\} \\
&= ||\psi||_\infty \sup_{g\in G}\{\int_H \beta(g^{-1}h)\ dm_H(h)\} \\
&= ||\psi||_\infty \\
&< \infty.
\end{aligned}$$

Also note that $T1 = 1$ by definition of β, and that T is a contraction by the second to last line. Furthermore, lemma 5.20 applies because β is a Bruhat function, this then tells us that $T\psi$ is continuous for every $\psi \in \mathcal{C}_b(H)$, in other words $T(\mathcal{C}_b(H)) \subseteq \mathcal{C}_b(G)$. Now, use the supposition that there exists a left invariant mean on $L^\infty(G)$ and theorem 5.5 to find a left invariant mean M on $\mathcal{C}_b(G)$. Then, the composition $M \circ T$ is a mean on $\mathcal{C}_b(H)$, that it is also left invariant is a consequence of the following, let $g \in G, h_0 \in H$ and $\psi \in \mathcal{C}_b(H)$, then

$$\begin{aligned}
T({}_{h_0}\psi)(g) &= \int_H \psi(h_0^{-1}h)\beta(g^{-1}h)\ dm_H(h) \\
&= \int_H \psi(h)\beta(g^{-1}h_0h)\ dm_H(h) \\
&= T\psi(h_0^{-1}g) \\
&=_{h_0} (T\psi)(g).
\end{aligned}$$

So, we have found a left invariant mean on $\mathcal{C}_b(H)$, theorem 5.5 then tells us that there exists one on $L^\infty(H)$ too, which is what we wanted to show. □

Corollary 5.23. *Let G be an amenable group (without topology endowed upon it), then every subgroup of G is amenable too.*

Proof. Endow G with the discrete topology, then there exists a left invariant mean on $L^\infty(G)$. By the subspace topology any subgroup H is also discrete and therefore closed in G. We can therefore apply theorem 5.22 which says that there exists a left invariant mean on $L^\infty(H)$, which is only possible if H is an amenable group. □

Corollary 5.24. *Any group containing the free group with two generators $\mathbb{F}_2$ as a subgroup is not amenable.*

Proof. By the previous corollary, this would imply that $\mathbb{F}_2$ is amenable, which is false since we showed that it is paradoxical in chapter 1 . □

Corollary 5.25. *The group $SO(3)$ is paradoxical.*

Proof. Let

$$a = \begin{pmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad b = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta \\ 0 & \sin\theta & \cos\theta \end{pmatrix}$$

with $\theta = \arccos(\frac{1}{3})$. Then, the subgroup generated by a and b is isomorphic to $\mathbb{F}_2$ (for details we refer to [5]), applying the previous corollary immediately tells us that $SO(3)$ is not amenable, in other words $SO(3)$ is paradoxical. □

6 References

[1] Bondy, J.A. and U.S.R. Murty (1976). *Graph Theory With Applications*. Fifth Printing (1982). Elsevier Science Publishing.

[2] Caspers M.P.T. (2016) *The geometry of discrete groups, Summerschool at Utrecht University*.

[3] Folland, G.B. (1994). *A Course in Abstract Harmonic Analysis*. First edition. CRC Press.

[4] Runde, V (2002). *Lectures on Amenability*. First edition. Springer-Verlag Berlin Heidelberg.

[5] Wagon, S (1985). *The Banach-Tarski Paradox*. Reprint (1993). Cambridge University Press.