Commutative von Neumann algebras and representations of normal Hilbert space operators, or "Spectral Measures without the Spectrum"

(title approved by Dr Leader)

Tim D. Austin

October 22, 2006

Introduction

What is a Hilbert space?

One can view a complex Hilbert space as the natural infinite dimensional analog of the spaces $\mathbb{C}^n$, and many of the most fundamental intuitions about Hilbert space geometry rest on a direct analogy with finite-dimensional geometry. While this intuition is certainly the right one for the basic foundations of the theory of Hilbert spaces, we should perhaps take ourselves to task about it: why should there be a sensible infinite-dimensional analog of the geometry of $\mathbb{C}^n$? Perhaps more importantly, why does such a thing matter?

In fact, Hilbert spaces are special precisely because of this direct generalization of finite dimensional geometry. Arguably they are the most artificial of all Banach spaces. While an M-space, for example, has a vast range of possible substructure, a Hilbert space has a geometric structure so homogeneous that even its own isometries cannot tell one point of the unit sphere from another. This absence of pathologies results in a highly agreeable geometry, and it is this itself, and the related behaviour of other constructions based on it, that makes Hilbert spaces worthy of study.

However, by itself the maxim that we should be able to extend finite-dimensional thinking will not get us very far. We need naturally-occurring Hilbert spaces to facilitate our study; and we find them in integration theory as $L^2_\mathbb{C}$ spaces (including, in particular, the space $\ell^2_\mathbb{C}$, when we consider counting measure on the natural numbers). (Note that we will persist in writing the subscript $\mathbb{C}$ in acknowledgement of the convention that L^2 is a real space.) In fact, we can say more: for many purposes, these are the *only* naturally-occurring Hilbert spaces that matter (an arguable exception to this is the Hardy space H^2, but in comparison this is of importance only under very special circumstances: the study of shift operators; see [11] for a comprehensive treatment).

A central message of this essay is that additional structures on a Hilbert space are often best understood by representing that Hilbert space itself in the form of one of these old friends. This can shed more light on the situation than a representation based only on higher-level structures such as algebras of operators. As Masamichi Takesaki put it in [15],

> "... However, an abstract Hilbert space alone cannot do much for us. We have to impose more structure on it. In fact, there is no abstract Hilbert space in eral life. Every Hilbert space arises through a specific construction. In various ways, the construction of Hilbert spaces involved integration... One has to reconstruct the entire structure through an appropriate way –..."

(It should be noted that here Takesaki was referring also to the study of noncommutative integration theory and its use in the study of arbitrary von Neumann algebras.)

It is worth noting that in adopting this philosophy we are implicitly recognizing the structures from which concrete Hilbert spaces can be defined – in particular, measure spaces and algebras – as being more fundamental than the Hilbert spaces themselves. We will therefore assume familiarity with a wide range of ideas from measure theory, although the most important will be reviewed briefly in Section 1.

The purpose of this essay

Some of the most satisfying theorems of mathematics are representation theorems, from which we see that every member of a class of abstractly-defined objects is actually a member of a family of naturally-occurring examples, up to an appropriate notion of isomorphism. Not only do such results have a great inherent beauty, but they also reduce the study of the abstract object to that of simpler, more fundamental structures, and so often spawn a large number of corollaries.

This is an essay about representing bounded normal operators on a Hilbert space. The best known result in this area is the Spectral Theorem, which dominates many introductory courses of operator theory (late undergraduate or beginning postgraduate). Regrettably, this theorem is often presented only as a natural corollary of the Gel'fand theory[1] of Banach algebras. While this latter theory is a masterpiece of twentieth century mathematics, an emphasis on this and the Spectral Theorem by themselves can be misleading. We shall develop a deeper and more insightful body of theory from which the Spectral Theorem follows as a natural corollary.

There is more than one fruitful approach to this task. Here we will go on to study commutative von Neumann algebras: commutative $*$-subalgebras of the algebra of all bounded operators on our Hilbert space that are closed for the weak-operator topology. Such an algebra is a commutative C*-algebra, and so the classic Gel'fand theory gives an essentially unique representation of such an algebra as the algebra of all continuous complex-valued functions on a compact Hausdorff space. One of the most natural routes to the results we wish to prove uses this representation theory by studying the properties of the compact Hausdorff space obtained in this way; however, this approach should not be the end of the story. It yields a representation of our von Neumann algebra as an algebra of functions, but tells us nothing about our original algebra of operators as an algebra *of operators*. More satisfactory is a representation of the underlying Hilbert space as a naturally-occurring Hilbert space which then induces an isomorphism from the original algebra of operators to an algebra of operators with a natural description in terms of the new Hilbert space (such an isomorphism of operator algebras given by an isomorphism of the underlying spaces is said to be spatial). Bearing in mind that our leading examples of Hilbert spaces are $L^2_{\mathbb{C}}$ spaces, a moment's thought suggests one such concrete algebra of operators in particular: the algebra of operators of the form M_h given by multiplication by a function h in $L^\infty_{\mathbb{C}}$.

In this essay we will see how to construct such a spatial isomorphism, and then how to derive the classical Spectral Theorem from it. This task may be approached in various ways. It is possible to use the Gel'fand representation theory as outlined above for a suitably-chosen von Neumann algebra $\mathcal{M}$ (and so appeal implicitly to all of the machinery involved in the study of Banach algebras and spectral theory) and then "read off" the desired spatial isomorphism. We will give a brief account of this approach in Section 4; a friendly in-depth account can be found in Conway's textbook [3], whereas a more cursory treatment, exemplifying many of the problems, as well as advantages, of a straightforward application of the Gel'fand theory, is in Rudin [13]. However, this is *not* an essay about spectral theory, and we will spend most of our time exploring a tougher but more elementary approach. In this we still start from a von Neumann algebra, but soon switch to the study of an associated resolution of the identity, and so link up with the theory of Boolean algebras. Once the relevant connections have been established, an appeal to the Stone Representation Theorem for Boolean algebras (in some ways analogous to the Gel'fand theory for C*-algebras, but much simpler) will lead us back to the same representation of $\mathcal{M}$ as given by the

[1] The phrase "Gel'fand theory" occurs frequently throughout this essay. In using it I do not mean to downplay the contributions of the other pioneers of Banach algebra theory, such as Naimark and Mazur; however, Gel'fand's towering contribution suggests this as a reasonable shorthand when referring to this overall corpus of ideas.

Gel'fand theory. In a sense this latter approach is made possible by the additional information about our von Neumann algebra furnished by its interaction with the underlying Hilbert space, which will allow us to use constructions not available when developing the fully abstract Gel'fand theory. The geometry of the underlying Hilbert space really does matter.

A routemap of these different approaches to the Spectral Theorem is shown in Figure 1. There are many other subtle variations that are not shown, but this outline should be enough to place any other auxiliary result or idea in its proper context.

None of this material is new; it gives rise to the body of results in the study of Hilbert space operators known as multiplicity theory, providing an essentially complete description of a special class of (not necessarily commutative) Hilbert space operator algebras, the type I von Neumann algebras.

1 Measure theory and Hilbert spaces

As has been mentioned, this essay is written with the conviction that the underlying objects of measure theory are more fundamental than complex Hilbert spaces (at least when these latter are infinite dimensional). This does not mean that they are necessarily simpler to understand than the definition of a complex Hilbert space, nor that they should be learnt first. However, since this essay is intended to show one example of how Hilbert spaces can often be best understood in terms of these more fundamental objects, this conviction is reflected in the priorities of the essay. Throughout this work we will assume familiarity with elementary measure theory (up to and including the basics of measure algebras and the theory of Stone spaces), functional analysis and Hilbert space geometry (including the Hilbert space Riesz Representation Theorem and the use of polarization identities). We will usually write (Ω, Σ, μ) for a measure space and $(\mathfrak{A}, \bar{\mu})$ for a measure algebra.

Given a measure space (Ω, Σ, μ), we define the space of **complex square-integrable functions** $\mathcal{L}^2_{\mathbb{C}}(\mu)$ as the space of measurable functions $f : \Omega_0 \to \mathbb{C}$ defined on a conegligible subset $\Omega_0 \subseteq \Omega$ and such that $\int_\Omega |f|^2 \, \mathrm{d}\mu < \infty$. This is a semi-Hilbert space with the semi-inner-product

$$\langle f, g\rangle = \int_\Omega \bar{f} g \, \mathrm{d}\mu,$$

and becomes the familiar Hilbert space $L^2_{\mathbb{C}}(\mu)$ upon quotienting by the equivalence relation that identifies functions on Ω that agree μ-a.e.

Similarly, we define the space $\mathcal{L}^\infty_{\mathbb{C}}(\mu)$ of **complex essentially bounded functions** as the space of functions $f : \Omega_0 \to \mathbb{C}$ defined on a conegligible subset $\Omega_0 \subseteq \Omega$ such that

$$\|f\|_\infty := \inf\{k > 0 : \ |f| \le k \ \mu - \text{a.e.}\} < \infty.$$

The Banach space $L^\infty_{\mathbb{C}}(\mu)$ is the quotient of this space by the equivalence relation of μ-a.e. equality, with the norm induced from the above expression.

Although $L^2_{\mathbb{C}}$ and $L^\infty_{\mathbb{C}}$ spaces are usually first introduced through the above construction, they are really associated to measure algebras rather than to specific measure spaces. Suppose we have a measure algebra $(\mathfrak{A}, \bar{\mu})$. There are various ways in which the classical function spaces of measure theory may be obtained directly from this. Perhaps the easiest is to let $S(\mathfrak{A})$ be the set of formal $\mathbb{C}$-linear combinations of members of $\mathfrak{A}$, and $S_0(\mathfrak{A})$ the subset of $\mathbb{C}$-linear combintions of members of $\mathfrak{A}$ of finite measure (following convention, when $a \in \mathfrak{A}$ is considered as a member of $S(\mathfrak{A})$ it will be written χ_a, to be thought of intuitively as the characteristic function of a).

Given $f = \sum_{i\le n} \lambda_i \chi_{a_i}$ and $g = \sum_{j \le m} \eta_j \chi_{b_j}$ in $S_0(\mathfrak{A})$ with $a_1, a_2, \ldots, a_n, b_1, b_2, \ldots, b_m$ finite-measure members of $\mathfrak{A}$, we define

$$\langle f, g\rangle := \sum_{i\le n, j\le m} \bar{\lambda}_i \eta_j \bar{\mu}(a_i b_j).$$

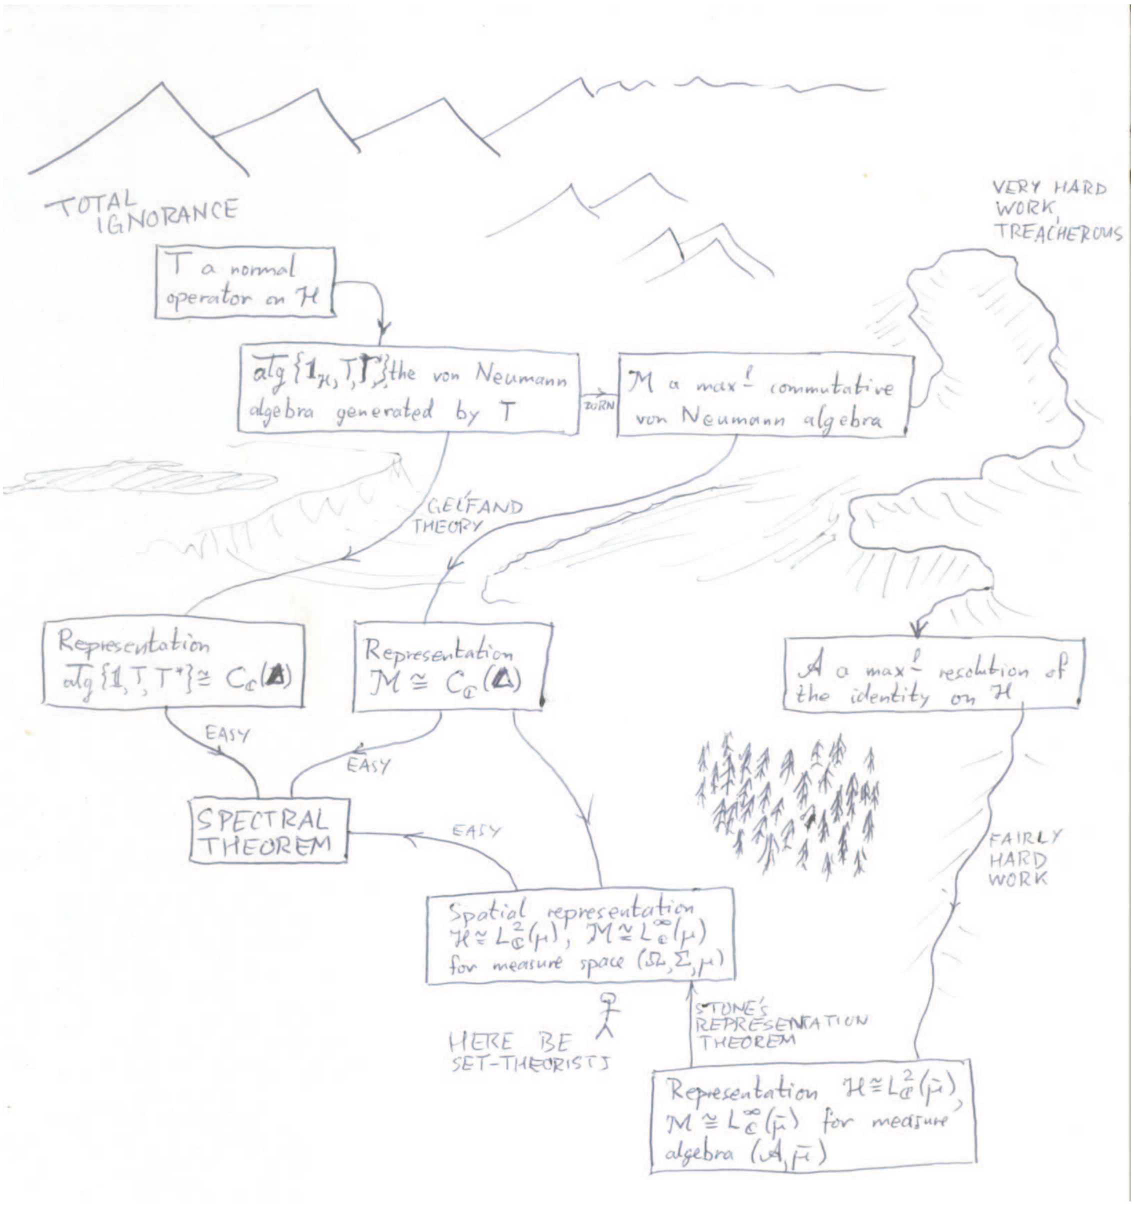

Figure 1: Routemap of essay

It is a routine exercise to show that this is an inner product on $S_0(\mathfrak{A})$. We may identify $L^2_{\mathbb{C}}(\mathfrak{A})$ as its completion in the induced norm.

Similarly, we define the norm $\|\cdot\|_\infty$ on $S(\mathfrak{A})$ by

$$\|f\|_\infty = \sup_{i\leq n} |\lambda_i|$$

when f is written as $\sum_{i\leq n} \lambda_i \chi_{a_i}$ with $a_1, a_2, \ldots, a_n$ non-zero in $\mathfrak{A}$ and disjoint, and identify $L^\infty_{\mathbb{C}}(\bar{\mu})$ as the completion of this space.

We will need to be able to move between these measure space and measure algebra notions of function space, and so will refer to the theory of Stone spaces and the resulting links between measure spaces and algebras. In particular, we will assume the following results.

Theorem 1.1 *Let $\mathfrak{A}$ be a Boolean algebra and let Z be the set of Boolean homomorphisms from $\mathfrak{A}$ to $\mathbb{Z}/2\mathbb{Z}$. The map $\mathfrak{A} \ni a \mapsto \hat{a} := \{\phi \in Z : \phi(a) = 1\} \subseteq Z$ is an isomorphism from $\mathfrak{A}$ onto a subalgebra of $\mathcal{P}Z$, and the topology on Z generated by $\{\hat{a} : a \in \mathfrak{A}\}$ is compact, Hausdorff and zero-dimensional with the sets $\hat{a}$ identified as precisely the clopen sets. Z is extremally-disconnected if and only if $\mathfrak{A}$ is Dedekind σ-complete.*

Theorem 1.2 *Let $(\mathfrak{A}, \bar{\mu})$ be a localizable measure algebra. Then it is isomorphic to the measure algebra of some localizable measure space (Ω, Σ, μ).*

A comprehensive treatment of this material can be found in [6].

2 Resolutions of the identity

As outlined in the introduction, we will want to construct a spatial isomorphism from an abstract Hilbert space $\mathcal{H}$ to an $L^2_{\mathbb{C}}$ space under which a given commutative von Neumann algebra is sent to the algebra of multiplication operators. We can hope that such an attempt will work only because of the rather special geometry of Hilbert space. Although there are many complex Banach spaces X for which some subalgebra of $\mathcal{B}(X)$ is (isometrically) spatially isomorphic to an algebra of multiplication operators on a Banach function algebra – any $L^p_{\mathbb{C}}$ space has this property, for example – we cannot usually hope to obtain so abstract a characterization of these operator algebras. The point about a Hilbert space is that it has no preferred way up: there is only one complex Hilbert space of each cardinal dimension (up to unitary isomorphism), and if a Hilbert space $\mathcal{H}$ is unitarily isomorphic to a space $L^2_{\mathbb{C}}(\mu)$ for some measure μ, then it will be unitarily isomorphic to it in many different ways, and so we can reasonably attempt to select one such isomorphism that suits our purposes once we know which operator algebra we are trying to represent.

Thus we will need a means of selecting an isomorphism from $\mathcal{H}$ to an $L^2_{\mathbb{C}}$ space. Considered a different way, this will show us how a measure algebra $(\mathfrak{A}, \bar{\mu})$ is reflected by additional structures defined on the associated "concrete" Hilbert space $L^2_{\mathbb{C}}(\bar{\mu})$. Ultimately we will want to choose this isomorphism so that it takes a particular operator algebra on $\mathcal{H}$ into a well-understood structure. We will see that this restricts our choice quite severely, but not completely; as one would expect, the more detailed the structure to be handled via the spatial isomorphism, the more constrained is the choice of this spatial isomorphism. We will identify just what restrictions we need to impose on this selection, and then will be able to make any further choices arbitrarily. The extent of this latter arbitrary choice can be described by specifying which automorphisms of the resulting concrete Hilbert space $L^2_{\mathbb{C}}$ preserve the "nice" representation of our original operator algebra.

As is often the case when assembling a proof of an existence theorem, the spatial isomorphism we need can be selected in various ways according to what kind of initial data we choose to work from. We will follow a standard approach, resting on the notion of a "resolution of the identity" (see Definition 2.1 below). In order to facilitate our intention to work more with measure algebras than with measure spaces for the time being, some of the definitions we use have been modified slightly from their usual form.

Recall first that if $\mathcal{H}$ is a complex Hilbert space then the collection of all its closed subspaces can be given the structure of a Dedekind complete lattice by defining

$$\bigvee_{i\in I} M_i := \overline{\bigoplus_{i\in I} M_i} = \overline{\mathrm{span}}\Big(\bigcup_{i\in I} M_i\Big)$$

and

$$\bigwedge_{i\in I} M_i := \bigcap_{i\in I} M_i$$

whenever $(M_i)_{i\in I}$ is a collection of closed subspaces of $\mathcal{H}$. This then translates into a Dedekind complete lattice structure on the family of all orthoprojectors on $\mathcal{H}$. Specifying to a commuting subfamily of orthoprojectors on $\mathcal{H}$, we have the following.

Definition 2.1 *Let $\mathcal{H}$ be a complex Hilbert space. A* ***resolution of the identity in*** $\mathcal{H}$ *is a commuting family of $\mathfrak{A}$ of orthoprojectors on $\mathcal{H}$ such that:*

1. *0 and $\mathbf{1}_{\mathcal{H}}$ are members of $\mathfrak{A}$;*
2. *if $E \in \mathfrak{A}$ then $\mathbf{1}_{\mathcal{H}} - E \in \mathfrak{A}$;*
3. *if $E, F \in \mathfrak{A}$ then $E \wedge F = EF \in \mathfrak{A}$ (this is still an orthoprojection because E, F commute);*
4. *if $(E_i)_{i\geq 0}$ is a sequence in $\mathfrak{A}$ then $\bigvee_{i\geq 0} E_i \in \mathfrak{A}$.*

It is now routine to check from the above conditions that if we define the complement of $E \in \mathfrak{A}$ to be $\mathbf{1}_{\mathcal{H}} - E \in \mathfrak{A}$ then a resolution of the identity $\mathfrak{A}$ is precisely a Boolean algebra of orthoprojectors on $\mathcal{H}$ that is sequentially order-closed as a sublattice of the lattice of all orthoprojectors on $\mathcal{H}$. We shall henceforth speak of $\mathfrak{A}$ as a Boolean algebra in this sense without further comment.

In a mild abuse of notation, we will sometimes also use the phrase "resolution of the identity" to refer to a sequentially order-continuous Boolean embedding of an arbitrary Dedekind σ-complete Boolean algebra $\mathfrak{A}$ into the family of orthoprojectors on $\mathcal{H}$ (this makes contact with the above definition by identifying $\mathfrak{A}$ with its image in the family of orthoprojectors).

We should not go any further before describing the leading example of a resolution of the identity and ultimate motivation for the above definition. Let $\mathcal{H} = L^2_{\mathbb{C}}(\bar{\mu})$ for a measure algebra $(\mathfrak{A}, \bar{\mu})$. For each $a \in \mathfrak{A}$ define the orthoprojector $E(a)$ on $\mathcal{H}$ as the multiplication operator $M_{\chi_a} : f \mapsto f \cdot \chi_a$; this projects onto the subspace of functions in $L^2_{\mathbb{C}}(\bar{\mu})$ that live on a. It is immediate to verify that E is a sequentially order-continuous Boolean embedding of $\mathfrak{A}$ into the lattice of orthoprojectors on $\mathcal{H}$, so $\{E(a) : a \in \mathfrak{A}\}$ is a resolution of the identity on $\mathcal{H}$. Ultimately we will seek to reverse this construction, by starting with a (maximal) resolution of the identity $\mathfrak{A}$ on an abstract Hilbert space $\mathcal{H}$ and using it to construct a measure algebra $(\mathfrak{A}, \bar{\mu})$ and an isomorphism from $\mathcal{H}$ to $L^2_{\mathbb{C}}(\bar{\mu})$. We will refer to a resolution of the identity of the above form as a **concrete** resolution of the identity.

Definition 2.2 *A resolution of the identity $\mathfrak{B}$ is a* ***refinement*** *of another resolution of the identity $\mathfrak{A}$ if $\mathfrak{A}$ is a subset of $\mathfrak{B}$.*

Definition 2.3 *Let $\mathfrak{A}$ be a resolution of the identity in $\mathcal{H}$. An* ***$\mathfrak{A}$-family*** *in $\mathcal{H}$ is a family $(x_E)_{E\in\mathfrak{A}}$ of vectors in $\mathcal{H}$ such that $Ex_F = x_E$ whenever $E \leq F$ in $\mathfrak{A}$. An $\mathfrak{A}$-family $(x_E)_{E\in A}$ is said to be* ***cyclic*** *if its linear span is dense in $\mathcal{H}$: $\overline{\mathrm{span}}\{x_E : E \in \mathfrak{A}\} = \mathcal{H}$. The families $(x_E)_{E\in\mathfrak{A}}$ and $(y_F)_{F\in\mathfrak{B}}$ are said to be* ***orthogonal*** *if $x_E \perp y_F$ for all $E \in \mathfrak{A}, F \in \mathfrak{B}$, or equivalently if $\overline{\mathrm{span}}\{x_E : E \in \mathfrak{A}\} \perp \overline{\mathrm{span}}\{y_F : F \in \mathfrak{B}\}$.*

More generally, if $\mathcal{I} \trianglelefteq \mathfrak{A}$ is an ideal in $\mathfrak{A}$ then an ***$\mathcal{I}$-family*** *in $\mathcal{H}$ is a family $(x_E)_{E\in\mathcal{I}}$ such that $Ex_F = x_E$ whenever $E \leq F \in \mathcal{I}$; cyclicity and independence for $\mathcal{I}$-families are defined analogously with the above.*

The simplest example of an $\mathfrak{A}$-family is given by $(Ex)_{E\in\mathfrak{A}}$ for a fixed $x \in \mathcal{H}$.

Definition 2.4 *A resolution of the identity $\mathfrak{A}$ on $\mathcal{H}$ is* ***maximal*** *if it has no non-trivial refinement.*

Lemma 2.5 *Let $\mathfrak{A}$ be resolution of the identity on $\mathcal{H}$. The following are equivalent:*

1. *$\mathfrak{A}$ is maximal;*
2. *any orthoprojector that commutes with every member of $\mathfrak{A}$ is itself a member of $\mathfrak{A}$.*

Furthermore if either of these conditions holds then there is a cyclic $\mathcal{I}$-family for some order dense ideal $\mathcal{I} \trianglelefteq \mathfrak{A}$.

It will follow from the work of this section that in fact the existence of such a cyclic family is equivalent to (1) and (2).

Proof (1 $\Leftrightarrow$ 2) Suppose $\mathfrak{A}$ is maximal and P is an orthoprojector which commutes with every $E \in \mathfrak{A}$. Letting $\mathfrak{B}$ be the family of orthoprojectors of the form $PE + (\mathbf{1}_{\mathcal{H}} - P)F$ for some $E, F \in A$, we see at once that $\mathfrak{B}$ is a refinement of $\mathfrak{A}$. Since $\mathfrak{A}$ is maximal we must have $\mathfrak{A} = \mathfrak{B}$, and so $P \in \mathfrak{A}$. The converse is immediate, for any non-trivial refinement of $\mathfrak{A}$ would contain orthoprojectors commuting with every member of $\mathfrak{A}$ but not themselves in $\mathfrak{A}$.

(2 $\Rightarrow$ 3) Suppose that (2) holds. For each $z \in \mathcal{H}$ let P_z be the projection onto the closed subspace $\overline{\operatorname{span}}\{Ez : E \in \mathfrak{A}\}$ of $\mathcal{H}$. Then clearly P_z commutes with every $E \in \mathfrak{A}$, so P_z itself is in $\mathfrak{A}$. A routine application of Zorn's lemma now gives a maximal family $Z \subseteq \mathcal{H} \setminus \{0\}$ with P_y, P_z orthogonal whenever $y, z \in Z$ are distinct; it follows that the sum $\sum_{z\in Z} P_z$ must be the whole of $\mathbf{1}_{\mathcal{H}}$, for otherwise we could choose some $y \in \mathcal{H}$ orthogonal to $\operatorname{img} P_z$ for every $z \in Z$, and so we could augment the family Z with y, contradicting maximality. Now we set

$$\mathcal{I} = \Big\{E \in \mathfrak{A} : E \le \sum_{z\in Z_0} P_z \text{ for some finite } Z_0 \subseteq Z\Big\},$$

the ideal in $\mathfrak{A}$ generated by the projections $\{P_z : z \in Z\}$. This is clearly order dense since $\sum_{z\in Z} P_z = \mathbf{1}_{\mathcal{H}}$. Defining

$$x_E := \sum_{z\in Z_0} Ez$$

whenever $E \le \sum_{z\in Z_0} P_z$, we obtain a cyclic $\mathcal{I}$-family $(x_E)_{E\in\mathcal{I}}$ for the order dense ideal $\mathcal{I}$, as required.

□

Lemma 2.6 *Suppose that $\mathfrak{A}$ is a maximal resolution of the identity on $\mathcal{H}$. Then the Boolean algeba $\mathfrak{A}$ is Dedekind complete and measurable.*

Proof Suppose that $S \subset A$, $S \ne \emptyset$. Then letting $P = \sup S$, we see that P is an orthoprojector that commutes with every $E \in \mathfrak{A}$, so it must be in $\mathfrak{A}$, and is therefore the desired supremum. To show that $\mathfrak{A}$ is measurable, it suffices to select an order dense ideal $\mathcal{I} \trianglelefteq \mathfrak{A}$ and a cyclic $\mathcal{I}$-family $(x_E)_{E\in\mathcal{I}}$ (this is possible by the second part of Lemma 2.5), and now define $\bar{\mu}$ on $\mathfrak{A}$ by

$$\bar{\mu}E = \sup\{\|x_F\|^2 : F \le E, F \in \mathcal{I}\};$$

it is routine to check that this is a localizable measure.

□

Definition 2.7 *Let $\mathfrak{A}$ be a resolution of the identity on $\mathcal{H}$. We say that $E \in \mathfrak{A}$ is* ***σ-finite*** *if there is some $x \in \mathcal{H}$ such that $Fx \ne 0$ for any $F \in \mathfrak{A}\setminus\{0\}$ with $F \le E$. We say that $\mathfrak{A}$ is* ***semi-σ-finite*** *if for any $E \in \mathfrak{A}$ there is some σ-finite $F \le E$, $F \ne 0$.*

I am not sure how natural the above definitions seem at first sight. They are motivated by the case of a concrete resolution of the identity: if $\mathcal{H} = L^2_{\mathbb{C}}(\bar{\mu})$ for a measure algebra $(\mathfrak{A}, \bar{\mu})$, then the orthoprojector $E(a)$ on $\mathcal{H}$ is σ-finite if and only if $a \in \mathfrak{A}$ is $\bar{\mu}$-σ-finite; this follows from Lemma 2.8 below. Is there a similar characterization of those $E(a)$ corresponding to an $a \in \mathfrak{A}$ of finite $\bar{\mu}$-measure? The answer is "no", as we will see after Corollary 2.9. This is an example of one of the subtle ways in which the Hilbert space $L^2_{\mathbb{C}}(\bar{\mu})$ and the concrete resolution of the identity $E : \mathfrak{A} \to \mathcal{B}(\mathcal{H})$ remember some, but not all, of the details of the measure algebra $(A, \bar{\mu})$ from which they were constructed.

Lemma 2.8 *A member $a \in A$ is of $\bar{\mu}$-σ-finite measure if and only if there is some $f \in L^2_{\mathbb{C}}(\bar{\mu})$ vanishing on no nonzero $b \leq a$ (that is, iff and only if the orthoprojector $E(a)$ is σ-finite).*

Proof Suppose first that $a = \bigvee_{i \geq 0} a_i$ with each a_i nonzero and of finite measure. Let $m_i = \bar{\mu} a_i$ for each $i \geq 0$ (so $m_i \neq 0$, since $a_i \neq 0$), and define f by

$$f = \sum_{i \geq 0} \frac{1}{2^i \sqrt{m_i}} \chi_{a_i}.$$

We can compute immediately that this f suffices.

On the other hand, if such an f exists then $\bar{\mu}[\![|f| > 2^{-i}]\!] < \infty$ for each $i \geq 0$, since $f \in L^2_{\mathbb{C}}(\bar{\mu})$, and $a = \bigvee_{i \geq 0}(a \cap [\![|f| > 2^{-i}]\!])$, so a is σ-finite.

□

Corollary 2.9 *Let A be a measurable Dedekind complete Boolean algebra and $\bar{\mu}$, $\bar{\nu}$ two localizable measures on it, and for $a \in A$ let $E_0(a)$ be the orthoprojector $f \mapsto f \cdot \chi_a$ in $L^2_{\mathbb{C}}(\bar{\mu})$ and $E_1(a)$ the same orthoprojector acting on $L^2_{\mathbb{C}}(\bar{\nu})$. Then the following are equivalent:*

1. *there is a unitary isomorphism $\Phi : L^2_{\mathbb{C}}(\bar{\mu}) \to L^2_{\mathbb{C}}(\bar{\nu})$ such that $\Phi \circ E_0(a) \circ \Phi^{-1} = E_1(a)$ for all $a \in A$;*
2. *there is a isomorphism $\Phi : L^2_{\mathbb{C}}(\bar{\mu}) \to L^2_{\mathbb{C}}(\bar{\nu})$ such that $\Phi \circ E_0(a) \circ \Phi^{-1} = E_1(a)$ for all $a \in A$;*
3. *a member $a \in A$ is $\bar{\mu}$-σ-finite if and only if it is $\bar{\nu}$-σ-finite.*

Proof $(1 \Rightarrow 2)$ This is clear.

$(2 \Rightarrow 3)$ Suppose first that such an isomorphism Φ exists, and let $a \in A$ be $\bar{\mu}$-σ-finite. Then be Lemma 2.8 there is some $f \in L^2_{\mathbb{C}}(\bar{\mu})$ with $E_0(b)f \neq 0$ for all $b \leq a$. It follows that

$$E_1(b) \circ \Phi(f) = \Phi \circ E_0(b)(f) \neq 0$$

for any $b \leq a$, since Φ is an isomorphism. Applying Lemma 2.8 again, this time with $\Phi(f)$ in place of f, we see that a must have $\bar{\nu}$-σ-finite measure. A symmetrical argument shows that any a with $\bar{\nu}$-σ-finite measure has $\bar{\mu}$-σ-finite measure

$(3 \Rightarrow 1)$ Applying Zorn's Lemma, we may choose a maximal family $(a_i)_{i \in I}$ of pairwise disjoint sets all of $\bar{\mu}$-σ-finite measure. Since $\bar{\mu}$ is semi-finite, such sets are order dense in the Boolean algebra A, and so $\bigvee_{i \in I} a_i = 1_A$ (where we have used the Dedekind completeness of A to deduce that the supremum on the left hand side exists). By condition (3), each a_i is also of $\bar{\nu}$-σ-finite measure. For each $i \in I$, the Radon-Nikodm Theorem now gives a non-negative measurable function h_i with $[\![h_i > 0]\!] \leq a_i$ and

$$\bar{\mu} b = \int_b h_i \, \mathrm{d}\bar{\nu}$$

for all $b \leq a_i$. We define the map $\Phi : L^2_{\mathbb{C}}(\bar{\mu}) \to L^2_{\mathbb{C}}(\bar{\nu})$ by linearity, continuity and the condition $\Phi(f) = f \cdot h_i$ whenever $[\![|f| > 0]\!] \leq a_i$. It is now straightforward to check that Φ has the desired properties.

□

It follows from this corollary that the measures $\bar{\mu}$ and $\bar{\nu}$ need not have the same finite-measure members of $\mathfrak{A}$ in order to give $L^2_{\mathbb{C}}$ spaces that are related by an isomorphism that preserves the concrete resolution of the identity $E : \mathfrak{A} \to \mathcal{B}(\mathcal{H})$. This tells us that the property that $a \in \mathfrak{A}$ be of finite measure (unlike the property of being of σ-finite measure) cannot be deduced from the concrete resolution of the identity E alone.

Having established Corollary 2.9, we can now use it to prove the following more general result.

Theorem 2.10 *Let $\mathfrak{A}$ be a maximal resolution of the identity on $\mathcal{H}$, and let $\bar{\mu}$ be a localizable measure on $\mathfrak{A}$. Then the following are equivalent:*

1. *there is a unitary isomorphism $\Phi : L^2_{\mathbb{C}}(\bar{\mu}) \to \mathcal{H}$ such that $\Phi^{-1} \circ E(a) \circ \Phi = M_{\chi_a}$ and $\|\Phi(\chi_a)\|^2 = \bar{\mu}a$ for all $a \in A$;*
2. *E and $\bar{\mu}$ have the same σ-finite members of A.*

These conditions characterize Φ up to a multiplicative unitary operator, in the sense that if Ψ is another such isomorphism then there is some measurable function $\phi : A \to \mathbb{T}$ such that

$$\Psi^{-1} \circ \Phi(f) = \phi \cdot f$$

for all $f \in L^2_{\mathbb{C}}(\bar{\mu})$.

Proof $(1 \Rightarrow 2)$ This is a restatement of Lemma 2.8.

$(2 \Rightarrow 1)$ By Lemma 2.5 there is an order-dense ideal $\mathcal{I} \trianglelefteq A$ with a cyclic $(E, \mathcal{I})$-family $(x_b)_{b \in \mathcal{I}}$. Now define $\bar{\nu}$ on A by

$$\bar{\nu}a = \sup\{\|x_b\|^2 : \ b \in \mathcal{I}, b \leq a\}.$$

It is clear that $\bar{\nu}$ is a measure on A (in particular, the cyclicity of $(x_b)_{b \in \mathcal{I}}$ implies that $\mathrm{img}\, E(c) \not\leq \overline{\mathrm{span}}\{x_b : \ b \in \mathcal{I}\}$ for every $c \in A \setminus \{0\}$, and so $x_b \neq 0$ for some $b \leq c$). Condition (2) now implies that $\bar{\mu}$ and $\bar{\nu}$ have the same σ-finite sets, and so by Corollary 2.9 it will suffice to construct a unitary isomorphism of the desired type from Φ from $L^2_{\mathbb{C}}(\bar{\nu})$ to $\mathcal{H}$. We construct our isomorphism Φ by first specifying it on step functions and then extending by continuity. Our hand is largely forced by the condition of matching up the members of E with the orthoprojectors given by multiplication by characteristic functions: we must set

$$\Phi\Big(\sum_{i \leq n} \lambda_i a_i\Big) = \sum_{i \leq n} \lambda_i x_{a_i}$$

for any $a_1, a_2, \ldots, a_n \in \mathcal{I}$. Since $\mathcal{I}$ is order dense and contains all the $\bar{\nu}$-finite members of A, such step functions are dense in $L^2_{\mathbb{C}}(\bar{\nu})$; the above map is clearly an isometry on the set of these step functions, and so extends to the desired unitary isomorphism, and we are done.

To prove the last statement of the theorem we need only note that if an automorphism T of $L^2_{\mathbb{C}}(\bar{\mu})$ commutes with all the operators M_{χ_a} then it must be equal to $M_{T(\chi_{1_A})}$ on the step functions, and hence on all functions by continuity, and if in addition T is an isometry then $T(\chi_{1_A})$ must take values in $\mathbb{T}$.

□

Theorem 2.11 *Let $\mathfrak{A}$ be a maximal resolution of the identity on $\mathcal{H}$. Then there are a localizable measure $\bar{\mu}$ on $\mathfrak{A}$ and a unitary isomorphism $\Phi : L^2_{\mathbb{C}}(\bar{\mu}) \to \mathcal{H}$ such that $\Phi^{-1} \circ E(a) \circ \Phi = M_{\chi_a}$ and $\|\Phi(\chi_a)\|^2 = \bar{\mu}a$ for all $a \in A$. These conditions characterize Φ up to a multiplicative unitary operator.*

Proof This follows upon combining Lemma 2.6 and Theorem 2.10.

□

The above thoerem gives us a way of selecting a measure algebra $(\mathfrak{A}, \bar{\mu})$ and an isomorphism $\Phi^{-1} : \mathcal{H} \to L^2_{\mathbb{C}}(\bar{\mu})$ under which a maximal resolution of the identity takes a nice form. This is the selection procedure we need in order to achieve our ultimate goal; in order to use it, we need now only reduce the problem of representing a von Neumann algebra to that of representing a maximal resolution of the identity.

3 Von Neumann algebras

Having mastered the technical machinery of the last section, we are ready to start on the core of the essay: manipulating a von Neumann algebra so that it can be understood in terms of that machinery.

3.1 Topologies on $\mathcal{B}(\mathcal{H})$, positivity and von Neumann algebras

Von Neumann algebras may be characterized in many ways, some abstract and others intrinsic to Hilbert space operator algebras. We shall work with an example of the latter type, but first we review some material on topologies and order structure on $\mathcal{B}(\mathcal{H})$.

Definition 3.1 *The* ***weak operator topology*** *on $\mathcal{B}(\mathcal{H})$ is the topology defined by the seminorms $T \mapsto |\langle Tx, y\rangle|$ as x, y range over $\mathcal{H}$; in this topology, a net of operators $(T_s)_{s\in S}$ converges to an operator T if and only if $\langle T_s x, y\rangle \to \langle Tx, y\rangle$ for every $x, y \in \mathcal{H}$.*

Definition 3.2 *Let $T \in \mathcal{B}(\mathcal{H})$. T is said to be* ***positive*** *(written $T \geq 0$) if it is self-adjoint (that is, $T = T^*$) and satisfies $\langle Tx, x\rangle \geq 0$ for all $x \in \mathcal{H}$. We will write $\mathcal{B}(\mathcal{H})^+$ for the set of positive operators in $\mathcal{B}(\mathcal{H})$. It is elementary to check that this is a closed convex cone in the norm topology of $\mathcal{B}(\mathcal{H})$. Since also $\mathcal{B}(\mathcal{H})^+ \cap (-\mathcal{B}(\mathcal{H})^+) = \{0\}$, this cone defines a partial order on $\mathcal{B}(\mathcal{H})$: we write $S \geq T$ if $S - T \geq 0$.*

The above ideas combine in the following standard result about order convergence which will be needed later on.

Lemma 3.3 *Suppose $(T_s)_{s\in S}$ is an upwards directed net of positive operators in $\mathcal{B}(\mathcal{H})$ (that is, S is an upwards directed poset, each T_s is positive and we have $T_{s_0} \geq T_{s_1}$ in $\mathcal{B}(\mathcal{H})$ whenever $s_0 \geq s_1$ in S). Suppose further that this net is bounded above: there is some $k \geq 0$ with $\|T_s\| \leq k$ for all $s \in S$. Then the net has a supremum T in $\mathcal{B}(\mathcal{H})$ with $\|T\| \leq k$, and this T is the limit of the net in the weak operator topology.*

Proof For each $x \in \mathcal{H}$, $(\langle T_s x, x\rangle)_{s\in S}$ is an upwards-directed net in $[0, \infty)$ bounded above by $k\|x\|^2$, so converges to its supremum $q(x)$, say. By the polarization identity it follows that for any $x, y \in \mathcal{H}$ the net $(\langle T_s x, y\rangle)_{s\in S}$ converges to a limit, say $\psi(x, y)$. It is clear that ψ is a sesquilinear form on $\mathcal{H}$ and that $|\psi(x, y)| \leq k\|x\|\|y\|$ for all $x, y \in \mathcal{H}$. As is standard, it follows from the Riesz Representation Theorem for Hilbert spaces that there is some $T \in \mathcal{B}(\mathcal{H})$ with $\langle Tx, y\rangle = \psi(x, y) = \lim_{s\uparrow} \langle T_s x, y\rangle$, so that $(T_s)_{s\in S}$ converges to T in the weak operator topology. It is now routine to check that T is actually the supremum of the family $\{T_s : s \in S\}$.

□

We are ready for the algebras.

Definition 3.4 *A* ***von Neumann algebra*** *is a $*$-subalgebra $\mathcal{M}$ of $\mathcal{B}(\mathcal{H})$ that is closed for the weak operator topology.*

We choose to work with this definition for two reasons: it is more in keeping with the emphasis on Hilbert space geometry to be maintained throughout, and it will make our subsequent work easier. It turns out that the above condition on $\mathcal{M}$ is equivalent to the requirement that $\mathcal{M}$ be equal to its own bicommutant. In addition, it is worth noting that it is possible to give an equivalent abstract definition, in which case the terminology **W*-algebra** is more frequently used. The following result is due to Sakai (see [14]):

Theorem 3.5 *For a* C**-algebra* $\mathcal{M}$ *the following are equivalent:*

1. *$\mathcal{M}$ is $*$-isomorphic to a von Neumann subalgebra of $\mathcal{B}(\mathcal{H})$ for some complex Hilbert $\mathcal{H}$;*
2. *$\mathcal{M}$ is Banach-space isomorphic to the dual of some other Banach space.*

A further characterization of W*-algebras in terms of their order structure and some additional technical assumptions is due to Kadison; see [7].

These abstract characterizations notwithstanding, W*-algebras occur regularly only in two settings: if they are commutative they arise as $L^\infty_{\mathbb{C}}$ spaces, and in general they arise as $*$-subalgebras of $\mathcal{B}(\mathcal{H})$. A vast body of theory for von Neumann algebras has been developed since the first seminal studies by Murray and von Neumann were published in the late 1930's and early 1940's ([8] – [10]), much of it motivated by the power of von Neumann algebras to provide a working formalism for quantum physics (see, for example, Bratteli and Robinson [1]). However, in this work we shall restrict ourselves to commutative von Neumann algebras, and so will not have recourse to most of the deeper theory.

The main task of this section is to obtain an isomorphism from a abstract Hilbert space $\mathcal{H}$ to an $L^2_{\mathbb{C}}$ space under which a maximal von Neumann algebra is sent to the algebra of multiplication operators. We will do this by showing that a maximal commutative von Neumann algebra contains a maximal resolution of the identity: loosely, this tells us that a "large" (and, in particular, maximal) commutative von Neumann algebra contains a "large" (in particular, maximal) collection of orthoprojectors. This will allow us to make contact with the results of Section 2, and so to select an isomorphism with the desired properties.

This is not easy if we are to avoid classical Banach algebra theory. Standard arguments to show that a commutative von Neumann algebra $\mathcal{M}$ contains a large supply of orthoprojectors (considered as the idempotents of the algebra) all rely on the Gel'fand representation of the algebra as the space of continuous functions $C(Z)$ for a suitable compact Hausdorff space Z; Dedekind completeness of $\mathcal{M}$ is then used to show that Z must be extremally disconnected and zero dimensional (a classical result of Stone), and so $C(Z)$ contains many idempotents in the form of characteristic functions of clopen sets.

However, we can take an alternative route relying only on (fairly) elementary Hilbert space geometry. We need to show that any member of a maximal von Neumann algebra can be approximated by finite positive linear combinations of orthoprojectors in the algebra (Corollary 3.11). Our rough plan for this proof reads as follows:

1. Show that any square of a self-adjoint operator has a unique positive square root;
2. Show that any self-adjoint operator S has a unique modulus $|S|$;
3. Show that any self-adjoint operator S can be split into positive and negative parts S_+, S_- with $S_+S_- = S_-S_+ = 0$;
4. Show that if $S \in \mathcal{M}$ is strictly positive then there is an orthoprojector $P \neq 0$ and an $\varepsilon > 0$ such that P commutes with every member of $\mathcal{M}$ and $S \geq \varepsilon P$;
5. Establish the approximation of arbitrary members of $\mathcal{M}$ by linear combinations of orthoprojectors, and hence the abundance of orthoprojectors.

It is not obvious at this stage just how each of the above steps follows from its predecessors. Some of these steps are quite technical; undoubtedly there are alternative itineraries that could be followed. Each step should become clear in due course.

We begin with the task of finding square roots of operators. In fact it is a standard result that every positive operator S on $\mathcal{H}$ has a unique positive square root, that is, a $T \geq 0$ with $T^2 = S$. The usual approach to this is via an alternative definition of positivity valid in an arbitrary C*-algebra, in terms of self-adjointness and the spectrum. The Gel'fand representation of a commutative C*-algebra then gives a quick proof of the existence of positive square roots for this latter definition of positivity, which may subsequently be shown to be equivalent to our notion of positivity in the case of Hilbert space operators. It is here, more than anywhere else, that we will feel the absence of the Gel'fand representation and the resulting powerful functional calculus. On this subject, it is worth noting the central role played in our proof by the power-series expansion for the function $t \mapsto \sqrt{1-t}$. The reader may recall that this series expansion is often used also in a proof of the Stone-Weierstrass theorem, which in turn is crucial in establishing the Gel'fand representation theorem. Although everything we do will be elementary, we are not so far from the Gel'fand approach after all.

Before proving the existence of square roots we need the following small lemma.

Lemma 3.6 *Suppose S is a self-adjoint operator on $\mathcal{H}$. Then*

$$\|S^2\| = \|S\|^2 = \sup_{\|x\|=1} |\langle S^2x, x\rangle|.$$

Proof The inequality $\|S^2\| \leq \|S\|^2$ follows from the sub-multiplicativity of the operator norm. Now suppose $x \in \mathcal{H}$ has $\|x\| = 1$. Then, using the fact that S is self-adjoint,

$$\|Sx\|^2 = \langle Sx, Sx\rangle = \langle S^2x, x\rangle \leq \|S^2x\|\|x\| = \|S^2x\|,$$

so $\|S\|^2 \leq \|S^2\|$. Thus all three quantities are actually equal.

□

This result is important because it tells us that whenever $T \geq 0$ has a self-adjoint square root, then we have $\|T\| = \sup_{\|x\|=1} |\langle Tx, x\rangle|$. In fact this result is true for any positive T, but this stronger result is difficult to prove without the continuous functional calculus, and we will manage without it.

We are ready to show the existence of square roots.

Lemma 3.7 *Suppose S is a self-adjoint operator on $\mathcal{H}$. Then there is a unique $T \geq 0$ in the uniform closure of* $\mathrm{alg}\{\mathbf{1}_{\mathcal{H}}, S\}$*, the uniform closure of the unital C*-algebra generated by S, that commutes with S and satisfies $T^2 = S^2$. This T also satisfies $T \geq \pm S$. Informally, "if an operator has a self-adjoint square root then it has a positive square root".*

Proof It clearly suffices to consider the case $\|S\| \leq 1$. The key idea of this proof is to use the Taylor expansion of the function $[0, 1] \ni t \mapsto \sqrt{1-t}$:

$$\sqrt{1-t} = 1 - \frac{1}{2}t - \frac{1}{2!}\frac{1}{2^2}t^2 - \frac{1}{3!}\frac{3}{2^3}t^3 - \cdots.$$

All the terms in this expansion after the first are negative, and so the series

$$\frac{1}{2}t + \frac{1}{2!}\frac{1}{2^2}t^2 + \frac{1}{3!}\frac{3}{2^3}t^3 + \cdots$$

converges absolutely for any $t \in [0, 1]$. This gives an expansion of $\sqrt{1-t}$, rather than of t, and so we need to write S^2 as $\mathbf{1}_{\mathcal{H}} - (\mathbf{1}_{\mathcal{H}} - S^2)$. Unfortunately, we do not yet know enough to be sure that the resulting series converges (in particular, we are not yet able to bound $\|\mathbf{1}_{\mathcal{H}} - S^2\|$), so we will need to digress slightly first. The proof proceeds in four steps.

STEP 1 First we show instead that $\mathbf{1}_{\mathcal{H}} - S^2$ has a positive square root, R say. For this we can apply the Taylor series expansion directly, defining

$$R := \mathbf{1}_{\mathcal{H}} - \frac{1}{2}S^2 - \frac{1}{2!}\frac{1}{2^2}S^4 - \frac{1}{3!}\frac{3}{2^3}S^6 - \cdots.$$

Since $\|S^2\| \leq 1$, we have $\|S^{2n}\| \leq \|S^2\|^n \leq 1$ for all $n \geq 1$, and so the absolute convergence of the power series for real $t \leq 1$ implies the convergence of the above sequence in the uniform operator topology (as the norms of the remainder terms are majorized by the remainder terms from the real case). This estimate also tells us that

$$\left\| \frac{1}{2}S^2 + \frac{1}{2!}\frac{1}{2^2}S^4 + \frac{1}{3!}\frac{3}{2^3}S^6 + \cdots \right\| \leq 1,$$

and so R is positive. Thus this R will suffice.

STEP 2 The importance of Step 1 is that we now know that the positive operator $\mathbf{1}_{\mathcal{H}} - S^2$ has a positive square root, and so as in the remark following Lemma 3.6 we see that

$$\|\mathbf{1}_{\mathcal{H}} - S^2\| = \sup_{\|x\|=1} |\langle x, x\rangle - \langle S^2x, x\rangle|.$$

Since $S^2 \geq 1$ and $\|S\| \leq 1$, this implies that $\|\mathbf{1}_{\mathcal{H}} - S^2\| \leq 1$, and hence that the power series that we really want,

$$T := \mathbf{1}_{\mathcal{H}} - \frac{1}{2}(\mathbf{1}_{\mathcal{H}} - S^2) - \frac{1}{2!}\frac{1}{2^2}(\mathbf{1}_{\mathcal{H}} - S^2)^2 - \frac{1}{3!}\frac{3}{2^3}(\mathbf{1}_{\mathcal{H}} - S^2)^3 - \cdots,$$

does converge uniformly. It therefore converges uniformly to a positive square root of S^2. Furthermore it is clearly in the uniform closure of $\operatorname{alg}\{\mathbf{1}_{\mathcal{H}}, S\}$ and so commutes with S.

STEP 3 Our next step is to show that $T \geq \pm S$. We know that $T^2 = S^2$ and $T \geq 0$ commutes with S. Hence $T^2 - S^2 = (T-S)(T+S) = 0$. Let P be the orthoprojector onto the kernel of $T-S$, so $Tx = Sx$ if and only if $x \in \operatorname{img} P$. From the above factorization we have $T(T+S) = S(T+S)$ and so $(T+S)x \in \operatorname{img} P$ for all $x \in \mathcal{H}$. Since $T+S$ is self-adjoint we have

$$(T+S)P = P(T+S)P = P(T+S)^*P = (P(T+S)P))^* = P(T+S),$$

so P commutes with $T+S$, similarly with $T-S$, and therefore also with T and S. Suppose $x \in \mathcal{H}$. If $x \in \operatorname{img} P$ then $x = Px$ and $Tx = Sx$, so

$$\langle Tx, x\rangle = \langle Sx, x\rangle,$$

implying $\langle Tx, x\rangle \geq \langle Sx, x\rangle$ and hence also, since $T \geq 0$, $\langle Tx, x\rangle \geq -\langle Sx, x\rangle$, so $\langle Tx, x\rangle \geq |\langle Sx, x\rangle|$. The result holds similarly for $x \in \operatorname{img}(\mathbf{1}_{\mathcal{H}} - P)$. Finally for arbitrary $x \in \mathcal{H}$ it follows that

$$\begin{aligned}
\langle Tx, x\rangle &= \langle T(Px + (\mathbf{1}_{\mathcal{H}} - P)x), x\rangle \\
&= \langle T(Px), Px\rangle + \langle T((\mathbf{1}_{\mathcal{H}} - P)x), (\mathbf{1}_{\mathcal{H}} - P)x\rangle \\
&\geq |\langle S(Px), Px\rangle| + |\langle S((\mathbf{1}_{\mathcal{H}} - P)x), (\mathbf{1}_{\mathcal{H}} - P)x\rangle| \\
&\geq |\langle S(Px), Px\rangle + \langle S((\mathbf{1}_{\mathcal{H}} - P)x), (\mathbf{1}_{\mathcal{H}} - P)x\rangle| = |\langle Sx, x\rangle|,
\end{aligned}$$

since P is idempotent and commutes with both T and S. Thus $T \geq \pm S$.

STEP 4 Finally we show uniqueness of T. Suppose $T_1, T_2 \geq 0$ and $T_1^2 = T_2^2$. The argument of Step 3 gives $T_1 \geq \pm T_2$ and $T_2 \geq \pm T_1$, so $T_1 = T_2$. This completes the proof.

□

The above is not the only approach that we can take to establishing the existence of square roots without appeal to the continuous functional calculus. Another possibility is to write T as an integral transform:

$$T := \frac{1}{\pi}\int_0^\infty \frac{1}{\sqrt{s}} S(s\mathbf{1}_\mathcal{H} + S)^{-1}\,\mathrm{d}s.$$

In this case we need to establish that the integral exists and has the desired properties; the full analysis is no easier than in the above approach, although it can be of use in the case of an abstract algebra (as opposed to an operator algebra). For more on this argument see Bratteli and Robinson [1], Theorem 2.2.10, and the references listed there.

Lemma 3.8 *Suppose $S \in \mathcal{B}_{\mathrm{sa}}(\mathcal{H})$. Then there are unique positive operators S_+, S_- in the uniform operator closure of* $\mathrm{alg}\{S, \mathbf{1}_\mathcal{H}\}$ *with $S = S_+ - S_-$ and $S_+S_- = S_-S_+ = 0$.*

Proof Define $|S|$ to be $\sqrt{S^2}$; by Lemma 3.7 this exists, is positive and commutes with S. Define

$$S_+ = \frac{1}{2}(|S| + S), \quad S_- = \frac{1}{2}(|S| - S).$$

We see at once that $S_+ + S_- = S$ and $S_+S_- = S_-S_+ = \frac{1}{4}(|S|^2 - S^2) = 0$.

On the other hand, these conditions on S_+, S_- give $S_+ + S_- \geq 0$ and $(S_+ + S_-)^2 = S_+^2 + S_-^2 = (S_+ - S_-)^2 = S^2$, so $S_+ + S_- = |S|$, by the uniqueness of the square root. Re-arranging now shows that S_+, S_- must take the above form in terms of $|S|$ and S, and so are unique.

□

We are now ready for the two pivotal lemmas. These will show us that we can approximate a self-adjoint member of a maximal commutative von Neumann algebra $\mathcal{M}$ by linear combinations of orthoprojectors in $\mathcal{M}$ in the weak operator topology. Loosely, this tells us that a "large" (and, in particular, maximal) commutative von Neumann algebra contains a "large" (in particular, maximal) collection of orthoprojectors, and so will allow us to make contact with the results of Section 2.

Lemma 3.9 *Suppose $\mathcal{M}$ is a commutative von Neumann algebra on $\mathcal{H}$, $T \in \mathcal{M}$ and $T \geq 0$, $T \neq 0$. Then there is a nonzero orthoprojector P that commutes with every member of $\mathcal{M}$ and an $\varepsilon > 0$ with $T \geq \varepsilon P$. If $\mathcal{M}$ is maximal it follows that $P \in \mathcal{M}$.*

Proof If there is any $\varepsilon > 0$ with $T \geq \varepsilon\mathbf{1}_\mathcal{H}$, then we need only set $P = \mathbf{1}_\mathcal{H}$, so henceforth we assume $T \ngeq \varepsilon\mathbf{1}_\mathcal{H}$ for every $\varepsilon > 0$. Since $T \geq 0$, $T \neq 0$, there is some $\varepsilon > 0$ and some $x_0 \in \mathcal{H}$ with $\|x_0\| \leq 1$ and $\langle Tx_0, x_0\rangle \geq 2\varepsilon$. Therefore $T - \varepsilon\mathbf{1}_\mathcal{H}$ is non-positive, and by the above property of x_0, it is also non-negative; hence

$$(T - \varepsilon\mathbf{1}_\mathcal{H})_+ \neq 0, \quad (T - \varepsilon\mathbf{1}_\mathcal{H})_- \neq 0.$$

Also, $(T - \varepsilon\mathbf{1}_\mathcal{H})_-(T - \varepsilon\mathbf{1}_\mathcal{H})_+ = 0$, so $\ker(T - \varepsilon\mathbf{1}_\mathcal{H})_-$ is a non-trivial proper subspace of $\mathcal{H}$. Let P be the orthoprojector onto this subspace, so P is nonzero. For any $x \in \mathcal{H}$ we have

$$\begin{aligned}
\langle Tx, x\rangle &= \langle(\varepsilon\mathbf{1}_\mathcal{H} + (T - \varepsilon\mathbf{1}_\mathcal{H})_+ - (T - \varepsilon\mathbf{1}_\mathcal{H})_-)x, x\rangle \\
&= \langle(\varepsilon P + (T - \varepsilon\mathbf{1}_\mathcal{H})_+)x, x\rangle + \langle(\varepsilon(\mathbf{1}_\mathcal{H} - P) - (T - \varepsilon\mathbf{1}_\mathcal{H})_-)x, x\rangle \\
&= \langle(\varepsilon P + (T - \varepsilon\mathbf{1}_\mathcal{H})_+)x, x\rangle + \langle(\varepsilon(\mathbf{1}_\mathcal{H} - P) - (T - \varepsilon\mathbf{1}_\mathcal{H})_-)(\mathbf{1}_\mathcal{H} - P)x, x\rangle \\
&= \langle(\varepsilon P + (T - \varepsilon\mathbf{1}_\mathcal{H})_+)x, x\rangle + \langle(\varepsilon\mathbf{1}_\mathcal{H} - (T - \varepsilon\mathbf{1}_\mathcal{H})_-)(\mathbf{1}_\mathcal{H} - P)x, (\mathbf{1}_\mathcal{H} - P)x\rangle,
\end{aligned}$$

for $(T - \varepsilon\mathbf{1}_\mathcal{H})_-(\mathbf{1}_\mathcal{H} - P) = (T - \varepsilon\mathbf{1}_\mathcal{H})_-$, by the definition of P. Now $\varepsilon P + (T - \varepsilon\mathbf{1}_\mathcal{H})_+ \geq \varepsilon P$ and $\varepsilon\mathbf{1}_\mathcal{H} \geq (T - \varepsilon\mathbf{1}_\mathcal{H})_-$, so the first term in the above sum is at least $\varepsilon\langle Px, x\rangle$ and the second is at least 0. Hence $T \geq \varepsilon P$.

Therefore this P will suffice if we can show that it commutes with any member of $\mathcal{M}$. This follows because $(T - \varepsilon \mathbf{1}_{\mathcal{H}})_- \in \mathcal{M}$ and so any $S \in \mathcal{M}$ must commute with $(T - \varepsilon \mathbf{1}_{\mathcal{H}})_-$. S will therefore map $\ker(T - \varepsilon \mathbf{1}_{\mathcal{H}})_-$ into itself. If S is self-adjoint it follows that $SP = PSP = (PSP)^* = (SP)^* = P^*S^* = PS$, and so S and P commute. Finally, since any member of $\mathcal{M}$ may be written as a linear combination of self-adjoint members of $\mathcal{M}$, the result follows. If $\mathcal{M}$ is maximal then it contains every operator that commutes with it, so $P \in \mathcal{M}$.

□

In fact a nonzero P with these properties necessarily exists in the closure of $\mathcal{M}$, even when $\mathcal{M}$ is not maximal, but the proof of this is a little trickier, and, once again, we do not need it. Like the other slight strengthenings of the auxiliary results we seen so far, it will follow once we have the Spectral Theorem.

Lemma 3.10 *Suppose $\mathcal{M}$ is a maximal commutative von Neumann algebra on $\mathcal{H}$, $T \in \mathcal{M}$ and $T \geq 0$. Then T is the supremum in $\mathcal{B}(\mathcal{H})$ of*

$$\{S : \ S \in \mathcal{M} \text{ is a finite positive linear combination of orthoprojectors in } \mathcal{M}, S \leq T\}$$

(referring to a linear combination as positive if all of its coefficients are positive).

Proof Let

$$A := \{S : \ S \in \mathcal{M} \text{ is a finite positive linear combination of orthoprojectors in } \mathcal{M}, S \leq T\}.$$

Clearly A is non-empty (it contains 0) and bounded above by T. If we can show that it is upwards directed, then it will follow by Theorem 3.3 that it has a supremum to which it converges in the weak operator topology.

To see that A is upwards-directed, suppose that $S = \sum_{i \leq n} \lambda_i E_i$ and $Q = \sum_{j \leq m} \eta_j F_j$ are in A with $E_1, E_2, \ldots, E_n$ mutually orthogonal and $F_1, F_2, \ldots, F_m$ mutually orthogonal (this latter condition may clearly always be imposed by splitting up terms in the sum if necessary). Since $A \subseteq \mathcal{M}$ it is commutative, so for each $i \leq n, j \leq m$ the product or orthoprojectors $E_i F_j = F_j E_i$ is itself an orthoprojector. We may therefore write

$$S = \sum_{i \leq n, j \leq m} \lambda_i (E_i F_j), \quad Q = \sum_{i \leq n, j \leq m} \eta_j (E_i F_j)$$

with the orthoprojectors $E_i F_j$, $i \leq n, j \leq m$, mutually orthogonal. Since S and T are both in A it follows that for each $i \leq n, j \leq m$ we have $S E_i F_j = \lambda_i E_i F_j \leq T E_i F_j$, and similarly $\eta_j E_i F_j \leq T E_i F_j$, so

$$S, Q \leq \sum_{i \leq n, j \leq m} \max\{\lambda_i, \eta_j\} E_i F_j \leq T \Big(\sum_{i \leq n, j \leq m} E_i F_j \Big) \leq T,$$

so S, Q have a common upper bound in A, and A is upwards-directed.

Let $R = \sup A$. By Theorem 3.3 this lies in the weak operator closure of A, and hence it is in $\mathcal{M}$. It follows also that it must satisfy $R \leq T$, by the weak-operator continuity of the map $S \mapsto \langle Sx, s \rangle$. Suppose $R \neq T$. Then $T - R \in \mathcal{M}$, $T - R \geq 0$, $T - R \neq 0$. Applying Lemma 3.9 we find some orthoprojection P in $\mathcal{M}$ and some $\varepsilon > 0$ with $R - T \geq \varepsilon P$. Now we have $S \leq R$ for every $S \in A$, and so $S + \varepsilon P \leq T$ for every $S \in A$. But this implies that $S + \varepsilon P \in A$ for every $S \in A$, and so $R \geq \sup\{S + \varepsilon P : \ S \in A\} = R + \varepsilon P$, a contradiction. Hence we must have $R = T$, and the result is proved.

□

Corollary 3.11 *Any member T of a maximal commutative von Neumann algebra $\mathcal{M}$ on $\mathcal{H}$ lies in the weak operator closure of the set of linear combinations of orthoprojectors in $\mathcal{M}$.*

Theorem 3.12 *Suppose $\mathcal{M}$ is a maximal commutative von Neumann algebra on $\mathcal{H}$ and $T \in \mathcal{B}(\mathcal{H})$ commutes with every orthoprojector in $\mathcal{M}$. Then T commutes with $\mathcal{M}$.*

Proof The map $\mathcal{M} \ni S \mapsto ST - TS \in \mathcal{B}(\mathcal{H})$ is continuous for the weak operator topology (since $\langle (ST-TS)x, y\rangle = \langle S(Tx), y\rangle - \langle Sx, T^*y\rangle$, a sum of continuous maps for the weak operator topology), and identically zero on the subset of linear combinations of orthoprojectors in $\mathcal{M}$, which is weak-operator-dense, by Corollary 3.11. Hence it must be zero on the whole of $\mathcal{M}$.

□

Lemma 3.13 *If $\mathcal{M}$ is a maximal commutative von Neumann algebra on $\mathcal{H}$ and $T \in \mathcal{B}(\mathcal{H})$ commutes with every member of $\mathcal{M}$ then $T \in \mathcal{M}$.*

Proof Suppose T is not in $\mathcal{M}$ and let $\mathcal{N}$ be the von Neumann algebra generated by T and $\mathcal{M}$ (that is, the intersection of all von Neumann algebras containing both). Now

$$\{S \in \mathcal{N} : \ S \text{ commutes with every member of } \mathcal{M} \text{ and with } T\}$$

is a von Neumann subalgebra of $\mathcal{N}$ that still contains T and $\mathcal{M}$, so it must be the whole of $\mathcal{N}$. Therefore T and any member of $\mathcal{M}$ lie in the centre of $\mathcal{N}$. This centre is therefore a von Neumann subalgebra of $\mathcal{N}$ that contains T and $\mathcal{M}$, so it must be the whole of $\mathcal{N}$, that is, $\mathcal{N}$ is commutative. But $\mathcal{N}$ properly contains $\mathcal{M}$, contradicting maximality.

□

Corollary 3.14 *Suppose $\mathcal{M}$ is a maximal commutative von Neumann algebra on $\mathcal{H}$. Then the set of orthoprojectors in $\mathcal{M}$ is a maximal resolution of the identity on $\mathcal{H}$.*

Proof Let $\mathfrak{A}$ be the set of orthoprojectors in $\mathcal{M}$. Since $\mathcal{M}$ is a commutative algebra, $\mathfrak{A}$ is a Boolean algebra, and since $\mathcal{M}_{\text{sa}}$ is sequentially-order complete, $\mathfrak{A}$ is closed under countable suprema. Therefore $\mathfrak{A}$ is a resolution of the identity. To show that $\mathfrak{A}$ is maximal we appeal to the second criterion of Lemma 2.5: suppose P is an orthoprojector commuting with every member of $\mathfrak{A}$. Then by Theorem 3.12 P commutes with every member of $\mathcal{M}$, and so by Lemma 3.13 must be in $\mathcal{M}$. Hence it is a member of $\mathfrak{A}$, and condition (2) of Lemma 2.5 is satisfied.

□

Now we are on the home straight. We set out to characterize a maximal commutative von Neumann algebra. This naturally gives rise to a Boolean algebra of orthoprojectors – the idempotents of the algebra – and the above Corollary tells us that this too is maximal. We may therefore appeal to the machinery laid out in Section 2. The final result is little more than a formality.

Theorem 3.15 *Suppose $\mathcal{M}$ is a maximal commutative von Neumann algebra on a complex Hilbert space $\mathcal{H}$. Then there is some measure algebra $(\mathfrak{A}, \bar{\mu})$ and a unitary isomorphism $\Phi : \mathcal{H} \to L^2_{\mathbb{C}}(\bar{\mu})$ under which $\mathcal{M}$ corresponds to the algebra of multiplication operators $\{M_h : \ h \in L^\infty_{\mathbb{C}}(\mathfrak{A})\}$.*

Proof Let $\mathfrak{A}$ be the algebra of orthoprojectors in $\mathcal{M}$. By Corollary 3.14, this is maximal, and so by Theorem 2.11 there is a unitary isomorphism $\Phi : \mathcal{H} \to L^2_{\mathbb{C}}(\bar{\mu})$ under which $\mathfrak{A}$ is sent to the concrete resolution of the identity. Since every member of $\mathcal{M}$ commutes with every member of $\mathfrak{A}$, the members of $\mathcal{M}$ must be sent to multiplication operators, and so, since the image of $\mathcal{M}$ contains the image of $\mathfrak{A}$, it must contain the weak-operator closure of its linear span. It follows that $\mathcal{M}$ must be sent precisely to $\{M_h : \ h \in L^\infty_{\mathbb{C}}(\bar{\mu})\}$.

□

Corollary 3.16 *Suppose $\mathcal{M}$ is a maximal commutative von Neumann algebra on a complex Hilbert space $\mathcal{H}$. Then there is some localizable measure space (Ω, Σ, μ) and a unitary isomorphism $\Phi : \mathcal{H} \to L^2_{\mathbb{C}}(\mu)$ under which $\mathcal{M}$ corresponds to the algebra of multiplication operators $\{M_h : \ h \in L^\infty_{\mathbb{C}}(\mu)\}$.*

Proof This follows at once from the above theorem, Theorem 1.2 and the correspondence between the function spaces of a measure space and a measure algebra.

□

Of course, neither of the above is quite the Spectral Theorem, which gives an explicit representation of one particular operator. In fact to make the transition to this theorem we do need to obtain the above Corollary first. The problem is that to find a suitable resolution of the identity to represent a single operator T (the so-called "Spectral Measure" of T) we need to be able to use expressions such as $[\![h \in A]\!]$ for $h \in L^\infty_{\mathbb{C}}$ and A a Borel subset of $\mathbb{C}$. This is difficult, although not impossible, when working purely in the measure algebra formulation: it requires some careful analysis to make the link between a member of $L^\infty_{\mathbb{C}}$ considered as a point in the completion of $S(\mathfrak{A})$ under the supremum norm and as something more like a function that can have a sensible notion of "image". Appealing to Stone's theorem and the above Corollary we can make light weather of this problem.

Corollary 3.17 (The Spectral Theorem) *Suppose T is a normal operator on a Hilbert space $\mathcal{H}$. Then there is some compactly-supported resolution of the identity E on the Borel σ-algebra of $\mathbb{C}$ such that*

$$T = \int_{\mathbb{C}} \lambda \, \mathrm{d}E(\lambda),$$

where this integral converges for the weak operator topology.

Proof It is elementary to check that the weak operator closure of $\mathrm{alg}\{\mathbf{1}_{\mathcal{H}}, T\}$ is a commutative von Neumann algebra, and a routine application of Zorn's Lemma allows us to extend this to a maximal commutative von Neumann algebra $\mathcal{M}$. By Corollary 3.16 there is a measure space (Ω, Σ, μ) such that $\mathcal{M}$ is spatially isomorphic to $L^\infty_{\mathbb{C}}(\mu)$ acting by multiplication on $L^2_{\mathbb{C}}(\mu)$. Suppose this spatial isomorphism sends T to the equivalence class of the essentially bounded function h. Define E by $E(A) = M_{\chi_{[\![h\in A]\!]}}$, noting that this is well-defined in terms of T because if h_1 and h_2 agree almost everywhere then $\mu([\![h_1 \in A]\!] \triangle [\![h_2 \in A]\!]) = 0$. We can now verify easily that for any square-integrable f on Ω we have

$$\Big(\int_{\mathbb{C}} \lambda \; M_{\chi_{[\![h\in \mathrm{d}\lambda]\!]}}\Big) f = h \cdot f$$

almost everywhere, and so

$$T = \int_{\mathbb{C}} \lambda \, \mathrm{d}E(\lambda).$$

Finally, E is compactly-supported because h is essentially bounded: $E(A) = \mathbf{1}_{\mathcal{H}}$ whenever $\mu[\![h \notin A]\!] = 0$.

□

In fact the last conclusion of the above theorem may be strengthened to give an intrinsic characterization of the support of E in $\mathbb{C}$. Let

$$A := \{z \in \mathbb{C} : \ \mu[\![h \in U]\!] > 0 \text{ for any neighbourhood } U \text{ of } z\}.$$

Since $\mathbb{C}$ is second-countable, we may cover the complement of A by a countable family $\mathcal{U}$ of open sets U with $\mu[\![h \in U]\!] = 0$, and so

$$\mu[\![h \notin A]\!] \leq \sum_{U \in \mathcal{U}} \mu[\![h \in U]\!] = 0.$$

Furthermore, we check easily that A is the smallest compact subset of $\mathbb{C}$ with this property, so it is the support of E. This set A is the set of points $z \in \mathbb{C}$ such that for no $\varepsilon > 0$ does the function h miss the open ball $B(z, \varepsilon)$ almost everywhere, and is therefore the set of those $z \in \mathbb{C}$ such that $h - z\chi_\Omega$ does not have an essentially bounded reciprocal. Translating this back into the language of abstract Hilbert space, this means that A is the set of those $z \in \mathbb{C}$ such that $T - z\mathbf{1}_{\mathcal{H}}$ has no inverse in $\mathcal{B}(\mathcal{H})$; this, of course, is precisely the spectrum of T.

4 Commutative von Neumann algebras and the Gel'fand theory

Some of the most powerful tools in the study of various classes of Banach algebra are the functional calculi, in particular the holomorphic, continuous and Borel functional calculi. These are often presented as a natural progression: the objects under study become increasingly specific, and the resulting functional calculi and associated representation theorems (representing subalgebras as function algebras) become more powerful. Certainly each of these functional calculi extends its predecessor and works only in a more restricted class of algebra. However, it is important to stress the different characters of the holomorphic functional calculus (and some of its other extensions such as the H^∞ functional calculus) and the continuous and Borel functional calculi. The former arises by considering a very special kind of algebra (the algebra of germs of analytic functions on some compact subset of $\mathbb{C}$) which we know can be 'fitted in' to any Banach algebra. This gives rise to a functional calculus which can be used under most circumstances. However, despite its power, this wide applicability carries the price that the holomorphic functional calculus usually misses much of the analytic structure in question. It gives a bare minimum of good algebraic behaviour which any Banach algebra must display, but cannot see any of a wide range of geometric or analytic properties.

The continuous functional calculus for a commutative C*-algebra and its specialization to the Borel functional calculus for a W*-algebra are quite different. Rather than trying to impose a structure that we understand onto a very generic object, these functional calculi arise as a by-product of the recognition of an abstractcommutative C*- or W*-algebra as a particular function algebra. This result is in many ways more delicate, requiring a detailed analysis of just how the different types of object involved interact. In return, it gives us an essentially complete description of commutative C*- and W*-algebras. The main result is the following.

Theorem 4.1 (The Gel'fand representation) *If $\mathcal{A}$ is a commutative C*-algebra and $\Delta \subset \mathcal{A}^*$ the set of $\mathbb{C}$-algebra homomorphisms $\mathfrak{A} \to \mathbb{C}$ with the restriction of the weak*-topology on $\mathcal{A}^*$, then Δ is compact Hausdorff and the Gelfand transform $\mathcal{A} \ni a \mapsto \hat{a} \in C_{\mathbb{C}}(\Delta)$ defined by $\hat{a}(\phi) = \phi(a)$ is a $*$-isomorhism.*

See Chapter 1 of [3]. Specializing slightly, we have the following.

Theorem 4.2 *If $\mathcal{M}$ is a commutative W*-algebra and $\Delta \subset \mathcal{A}^*$ the set of $\mathbb{C}$-algebra homomorphisms $\mathfrak{A} \to \mathbb{C}$ with the restriction of the weak*-topology on $\mathcal{A}^*$, then Δ is compact, Hausdorff, zero-dimensional and extremally disconnected.*

See [14]. In the case of a maximal commutative von Neumann algebra $\mathcal{M}$ acting on a Hilbert space $\mathcal{H}$, we may use the above result to construct a localizable Borel measure μ on Δ such that the action of $\mathcal{M}$ on $\mathcal{H}$ is spatially isomorphic to the action of $L^\infty_{\mathbb{C}}(\mu)$ (which can be identified with $C_{\mathbb{C}}(\Delta)$ owing to the special form of Δ) on $L^2_{\mathbb{C}}(\mu)$. We thus return to Corollary 3.16; of course, in fact the space Δ constructed here is just the same as the Stone space Z of the Boolean algebra $\mathfrak{A}$ needed to deduce Corollary 3.16 from Theorem 3.15.

Epilogue

The results of this paper are not substantially new; indeed, they form a small part of "multiplicity theory", which yields a complete classification of commutative von Neumann algebras. In the case of a separable Hilbert space $\mathcal{H}$ this takes the following form.

Theorem 4.3 *Suppose $\mathcal{M}$ is a commutative von Neumann algebra acting on $\mathcal{H}$. Then there are compact metric space $X_\infty, X_1, X_2, \ldots$ and Radon measures $\mu_\infty, \mu_1, \mu_2, \ldots$ on them such that $\mathcal{M}$ acting on $\mathcal{H}$ is spatially isomorphic to $L^\infty_{\mathbb{C}}(\mu_\infty)^\infty \oplus L^\infty_{\mathbb{C}}(\mu_1) \oplus L^\infty_{\mathbb{C}}(\mu_2)^2 \oplus \cdots$ (ℓ^∞ direct sum) acting by multiplication on $L^2_{\mathbb{C}}(\mu_\infty)^\infty \oplus L^2_{\mathbb{C}}(\mu_1) \oplus L^2_{\mathbb{C}}(\mu_2)^2 \oplus \cdots$ (Hilbert space direct sum).*

See Chapter 7 of [4], or the very detailed treatment in [5]. Fully general multiplicity theory in the case of a non-separable Hilbert space is more difficult, relying on a family of more advanced measure-theoretic ideas developed by Halmos for the purpose; see the expository article [2].

There are various other related questions that we have not so far even considered. The theory of (not necessarily commutative) von Neumann algebras is vast and beautiful, requiring additional ways of thinking which do not come up at all in the results explained above. Two classic treatments of the groundwork are those of Dixmier [5] and Sakai [14].

We have also not mentioned the spectral theorem for an unbounded normal operator. Suffice it to say that this is usually proved by appealing to the bounded case and being cunning, but the required build-up through the basic structure theory of unbounded operators would take us too far off-course. Once again, Conway's book [3] provides a readable account.

I have not found any references for a derivation of the spectral theorem for bounded normal operators entirely independently of the Gel'fand theory, and so this approach had to be developed for this essay.

References

[1] Bratteli O. & Robinson D.W., *Operator Algebras and Quantum Statistical Mechanics*, vols. 1 & 2, Springer-Verlag, Berlin, 1987;

[2] Brown A., "A version of multiplicity theory", *Mathematical Surveys no. 13: Topics in Operator Theory*, American Mathematical Society, Providence, 1974;

[3] Conway J.B., *A Course in Functional Analysis*, Springer-Verlag, New York, 1990;

[4] Conway J.B., *A Course in Operator Theory*, American Mathematical Society, Providence, 2000;

[5] Dixmier J., *Von Neumann Algebras*, North-Holland, Amsterdam, 1981;

[6] Fremlin D.H., *Measure Theory*, "Volume 3: Measure Algebras", Torres Fremlin, Colchester, 2002;

[7] Kadison R.V, "Operator algebras with a faithful weakly closed representation", *Ann. Math.* 64 (1956), 175 – 181;

[8] Murray F.J. & von Neumann J., "On rings of operators", *Ann. Math.* 37 (1936), 116 – 229;

[9] Murray F.J. & von Neumann J., "On rings of operators II", *Trans. Amer. Math. Soc.* 41 (1937), 208 – 248;

[10] Murray F.J. & von Neumann J., "On rings of operators IV", *Ann. Math.* 44 (1943), 716 – 808;

[11] Nikol'skii N.K., *Treatise on the Shift Operator*, Springer-Verlag, Berlin, 1986;

[12] Pedersen G., *C^*-algebras and their Automorphism Groups*, Academic Press, London, 1979;

[13] Rudin W., *Functional Analysis*, McGraw-Hill, New York, 1979;

[14] Sakai S., C^**-Algebras and* W^**-Algebras*, Springer-Verlag, Berlin, 1998;

[15] Takesaki M., *Structure of Factors and Automorphism Groups*, American Mathematical Society, Providence, 1983.

Tim D. Austin
Trinity College
Cambridge
September 2004