

Contextuality and Noncommutative Geometry in Quantum Mechanics



Nadish de Silva

Merton College

University of Oxford

A thesis submitted for the degree of

DPhil in Computer Science

For Ammi and Thaththi.

Abstract

It is argued that the geometric dual of a noncommutative operator algebra represents a notion of *quantum state space* which differs from existing notions by representing observables as maps from states to outcomes rather than from states to distributions on outcomes. A program of solving for an explicitly geometric manifestation of quantum state space by adapting the spectral presheaf, a construction meant to analyze contextuality in quantum mechanics, to derive simple reconstructions of noncommutative topological tools from their topological prototypes is presented.

We associate to each unital C^* -algebra \mathcal{A} a geometric object—a diagram of topological spaces representing quotient spaces of the noncommutative space underlying \mathcal{A} —meant to serve the role of a *generalized Gel'fand spectrum*. After showing that any functor F from compact Hausdorff spaces to a suitable target category \mathcal{C} can be applied directly to these geometric objects to automatically yield an extension \tilde{F} which acts on all unital C^* -algebras, we compare a novel formulation of the operator K_0 functor to the extension \tilde{K} of the topological K -functor. We then conjecture that the extension of the functor assigning a topological space its topological lattice assigns a unital C^* -algebra the topological lattice of its primary ideal spectrum and prove the von Neumann algebraic analogue of this conjecture.

Contents

0	Acknowledgements	6
1	Introduction	7
2	Background and Motivation	11
2.1	Mathematical Models of Physical Systems	12
2.2	State-observable Duality	14
2.3	Quantum State Space	17
2.4	The Noncommutative Geometry of C^* -algebras	19
2.5	The ‘Geometry’ in Noncommutative Geometry	21
3	Contextuality	25
3.1	Ontology and Epistemology	26
3.1.1	Ontological models	27
3.2	Contextuality in quantum theory: the Kochen-Specker Theorem . . .	30
3.3	The spectral presheaf	32

3.3.1	Points: the Kochen-Specker theorem	33
3.3.2	Distributions: Gleason's theorem	35
4	Spatial Diagrams	38
4.1	The Categories of all Diagrams in \mathcal{C}	39
4.2	Semispectral Functors	40
4.3	Spatial diagrams	42
4.4	Example: $M_2(\mathbb{C})$	44
5	Extensions of Topological Functors	46
5.1	The Generalized Limit and Colimit Functors	47
5.2	Extensions of Functors	50
5.3	Theorems of quantum foundations	52
6	Reconstructing Operator K-theory	53
6.1	Topological K -theory	54
6.2	Operator K -theory	56
6.2.1	Unitalization	58
6.2.2	Stability	59
6.3	Alternative definition of operator K_0 functor	62
6.4	The extension of topological K -theory	64

7	Noncommutative Topology	68
7.1	The primitive ideal space	69
7.1.1	The spectrum of commutative rings	69
7.1.2	The primitive ideal space	70
7.2	Von Neumann Algebras	71
7.2.1	Projections	71
7.2.2	Ideals	73
7.3	Total and partial ideals	75
7.3.1	Motivation	75
7.3.2	Partial and total ideals of C^* -algebras	77
7.3.3	Partial and total ideals of von Neumann algebras	79
7.4	Technical Preliminaries	82
7.4.1	Little lemmata	82
7.4.2	Partial orthogonality	83
7.4.3	Main lemma	84
7.5	Main theorem	87
7.6	Conclusions	89
8	Conclusions	90

Chapter 0

Acknowledgements

It is a pleasure to thank my supervisors, Samson Abramsky and Bob Coecke, as well as Chris Heunen, Andreas Döring, and Jamie Vicary for their guidance and encouragement during this project. I also wish to thank Kobi Kremnitzer, George Elliott, Rui Soares Barbosa, and Brent Pym for many crucial, constructive mathematical conversations. I would like to acknowledge the financial support of Merton College, the Oxford University Computing Laboratory, the Clarendon Fund, and the Natural Sciences and Engineering Research Council of Canada.

Chapter 1

Introduction

The discovery of the quantum effected a revolutionary shift in our fundamental conceptions of reality so profound that it necessitated an entirely new mathematics with which to frame physical theories. A century later, elucidating the foundational implications of those ideas remains a highly active area of both mathematics and physics.

Von Neumann provided a common mathematical structure which unified various proposed formalisms for quantum theory [70]. In his approach, algebraic relations between observable quantities are emphasized and the essential departure from the framework of classical theory is the presence of “noncommuting” quantities, i.e. properties of a system which cannot be jointly measured. Despite the apparently completely novel nature of this abstract, algebraic (commutator) formalism, there lies, beneath the surface, a great deal of structural similarity with classical mathematics and the geometric structures used to model classical physics. Indeed, modern mathematicians consider von Neumann algebras to be a generalization of classical probability spaces. C^* -algebras (which are also employed to model quantum systems [66]) are considered to be generalizations of classical topological spaces.

Although geometric intuition played a key role in formulating classical mechanics and early quantum theory, as well as modern mathematical studies of operator algebras, explicitly geometric constructions associated to noncommutative algebras are generally eschewed in favour of purely algebraic treatments which obscure similarities with classical theory. In this dissertation, we attempt to define an explicitly geometric interpretation for a noncommutative algebra of observables. This is done by accounting for contextuality: the phenomenon of the outcome of a measurement depending on the procedure used (or more formally, the choice of simultaneously measured quantities).

The germ of this idea was contained in Isham and Butterfield’s reformulation (with Hamilton) of the Kochen-Specker theorem in terms of the lack of a global section of the spectral presheaf associated with a quantum system [34]. That original insight has sparked research in a number of directions (for example: the work of Isham-Döring et al. [19, 20, 21, 22] and Heunen et al. [37, 38, 39] on developing an alternative formalism for quantum mechanics in terms of presheaves over contexts; Abramsky-Brandenburger et al. [1] on utilizing presheaves and sheaves as a framework for abstractly studying contextuality in theories more general than quantum mechanics).

With this thesis we argue that the passage from classical to quantum descriptions of the core operational content of physical theories, characterized mathematically by the transition from a commutative algebra of observables to a noncommutative one, can also be understood geometrically as a shift from classical sample spaces to context-indexed diagrams of spaces. Second, we demonstrate how, by considering other sorts of diagrams of spaces, we gain a new approach to the study of the noncommutative extension of topology.

In the background section, we describe how the state-observable dualities in mathematical models of physics are to be understood as examples of the more general phenomenon of the geometry-algebra dualities which arise throughout mathematics and logic. The particular example of Gel’fand duality is examined closely, as it models state-observable duality in classical mechanics. It is argued that a geometric dual for noncommutative C^* -algebras would represent a formal analogue of pure state space for quantum mechanics. We describe our approach of finding such a geometric dual by solving for the correct geometric structure which yields reconstructions of noncommutative geometric tools.

Chapter 2 contains an exposition of the Kochen-Specker theorem and of the phenomenon of contextuality in quantum theory. We provide a self-contained, ele-

mentary account of the spectral presheaf as a context-indexed collection of sample spaces. It is shown how, considering the spectral presheaf to represent to a notion of geometric space, the Kochen-Specker theorem can be seen as its lack of global points whereas Gleason's theorem can be understood as a correspondence between quantum states and probability distributions on the space [17].

The final chapters contain our primary original technical contributions. We introduce a new approach to studying the noncommutative geometry of C^* -algebras and to investigating possible generalizations of the Gel'fand spectrum to noncommutative algebras. This is accomplished by defining a simple method for naturally extending functors which act on topological spaces to ones acting on all C^* -algebras which follows from associating diagrams of spaces to noncommutative algebras. These diagrams are interpreted as keeping track of all the tractable quotient spaces of the noncommutative space represented by an algebra.

It is shown that functorially associating a diagram of compact Hausdorff spaces (meant to generalize the notion of spectrum) to a noncommutative algebra leads to a simple method of extending topological functors. In particular, a new definition of operator K -theory given in terms of topological K -theory is presented. A conjecture about using the same method to extend the notion of an open set to the corresponding notion of closed, two-sided ideal is made; the von Neumann algebraic version of this conjecture is proved.

Chapter 2

Background and Motivation

The notion of state and observable of an (operational) physical theory is defined. State-observable duality is described and some examples of geometric-algebraic dualities are presented. It is argued that the geometric dual of a noncommutative operator algebra represents a notion of *quantum state space* which differs from existing notions by representing observables as functions from states to outcomes rather than from states to distributions on outcomes. Noncommutative geometry, the algebraic study of the “quantum geometries” underlying noncommutative algebras is presented. A program of finding an explicitly geometric manifestation of quantum state space by adapting structures related to the phenomenon of contextuality in quantum mechanics to derive simple reconstructions of noncommutative topological tools from their topological prototypes is presented.

2.1 Mathematical Models of Physical Systems

A physical theory provides mathematical models for the physical systems which are in the domain of applicability of the theory. These models consist of some mathematical structures (i.e. manifolds, vector spaces, etc.) and prescriptions for how they can be used to compute numerical predictions of experimental procedures. It may also provide a description of dynamics—how systems evolve with time—but we ignore this aspect for now.

The empirical data which a successful model must predict can be summarily expressed as a map from experiments to predictions. An experiment is simply a pair consisting of a method for preparing a system and a procedure for measuring some property. Thus, a description of this data must specify two primitive collections. The first is the collection P of preparations: all the possible ways of manipulating a system to ready it for an experiment. The second is the collection M of measurements: all the possible procedures for probing such a prepared system in order to witness an experimental event. (One may choose to treat the measurable space (O, E) of possible outcomes of measurements as another variable.) The predictive content of the theory is then completed by specifying the collection F of probability distributions $f(e|p, m)$ of observing an event $e \in O$ upon performing the measurement procedure m on a system prepared by a procedure p .

Definition 2.1.1. *An operational theory is a triple (P, M, F) where P is a set of preparations, M a set of measurement procedures, and a set F of probability distributions $f(-|p, m)$, indexed by pairs $(p, m) \in P \times M$, on the space of outcomes. [51, 52]*

A description of a physical system at this operational level will be rich with redundancies. For example, trivial differences in measurement procedures (e.g. using apparatuses facing different cardinal directions) will not affect the predicted outcomes for any preparation of the system. In this situation, we have an intuitive sense

that there is one single property of the system which is being commonly measured by both these procedures. Two procedures which extract the same fundamental property should yield identical data for all experiments and, so, this sort of redundancy is formalized by considering operational equivalence classes [9, 64, 66].

Definition 2.1.2. *The observables of an operational theory (P, M, F) are the equivalence classes of measurements under the equivalence relation*

$$m_1 \sim m_2 \iff \forall p \in P, f(-|p, m_1) = f(-|p, m_2)$$

Similarly, two methods of preparing a system may leave it in physical configurations which are essentially the same. A necessary condition for two physical configurations being the same is that there exists no empirical means of distinguishing them.

Definition 2.1.3. *The states of an operational theory (P, M, F) are the equivalence classes of preparations under the equivalence relation*

$$p_1 \sim p_2 \iff \forall m \in M, f(-|p_1, m) = f(-|p_2, m)$$

We thus think of states and observables as being closer to the more ontologically fundamental concepts of *real* configurations and properties of a system but we are not justified in strictly identifying these pairs of concepts. In this extensional viewpoint, states can be identified with their (distributions on) outcomes for all observables and observables can be identified with their (distributions on) outcomes for all states. This fundamental symmetry leads to the notion of state-observable duality; a theme which recurs throughout mathematics and logic, manifesting as categorical dualities between geometry and algebra; semantics and syntax.

2.2 State-observable Duality

A good mathematical model should contain representations of states and observables rather than preparations and measurements. The collections of states and observables are often endowed with additional mathematical structure. The observables, being representatives of quantities which vary with state, are generally endowed with algebraic structure reflecting the arithmetic of quantities. The states, on the other hand, are endowed with geometric structure. Intuitively, states are close to each other when they represent configurations of a system which share similar physical properties as measured by experiments.

Important examples are those classical systems which can be modeled in terms of Poisson geometry [53]. The collection of pure states is in fact a geometric space: a Poisson manifold. This justifies use of the terminology *state space*. Any smooth real-valued map from this manifold can be taken to represent an observable quantity and taken together, these maps form a commutative algebra with pointwise operations. In this case, the Poisson bracket provides the additional structure of a Lie algebra. Hence, we refer to the *algebra of observables*.

In the above example, predictions for the outcomes of experiments are deterministic and observables are explicitly represented as quantity-valued functions on the state space. However, the fact that a pairing of a state with an observable results in a quantity means that fixing a state yields a quantity-valued function on the collection of observables. Identifying a state with the function on observables it defines allows realizing the state space as a space of functions from the algebra of observables to an algebra of quantities.

This perspective is common in duality theory. The simplest example is the Stone-type duality between the categories Set of sets and functions and caBa of complete, atomic Boolean algebras and complete Boolean algebra homomorphisms [65, 42].

A functor maps a set S to the Boolean algebra $\text{Hom}_{\text{Set}}(S, 2)$ of functions to $2 = \{0, 1\}$ and a function $f : S \rightarrow T$ to a caBa-morphism via pullback: $f^*(g) = g \circ f$. We can use, in the opposite direction, the functor $\text{Hom}_{\text{caBa}}(-, 2)$, where 2 is the two element Boolean algebra, to complete the duality of these categories. This is an equivalence between a category of geometric objects (sets can be seen as trivial geometries with no structure beyond cardinality) and algebraic objects.

An duality of the same form (defined by Hom functors to a dualizing object 2) exists between the categories of Stone spaces and Boolean algebras. The geometric nature of Stone spaces, which are certain kinds of topological spaces, is clearer in this instance. This example also clearly demonstrates a logical form of duality between semantics and syntax. The algebraic category of Boolean algebras can be seen as the category of propositional theories whereas the geometric category of Stone spaces is the category of corresponding spaces of two-valued models [27].

A classic example of geometric-algebraic duality, which informs the work in Chapter 7, is that which exists between commutative, unital rings and affine schemes [35]. Given such a ring R , one can define a topological space $\text{Spec}R$ called the *prime spectrum* (or, just *spectrum*) whose points are the prime ideals of R and whose open sets are indexed by ideals of R . One can then define a sheaf of commutative rings on $\text{Spec}R$ such that the stalk at a prime ideal p is the localization of R at p , turning $\text{Spec}R$ into a locally ringed space. The locally ringed spaces which arise in this way are called *affine schemes*. The commutative ring giving rise to an affine scheme can be recovered by taking the ring of global sections of the scheme. In this way, a geometric dual to the category of commutative, unital rings is constructed and geometric tools and reasoning can be brought to bear in subjects which make use of commutative rings, such as number theory.

The most important example for our purposes is the Gel'fand duality between the category KHaus of compact, Hausdorff spaces and continuous functions and

the category \mathbf{uCC}^* of unital, commutative C^* -algebras and $*$ -homomorphisms [29]. Under this duality, a space X corresponds to the commutative unital C^* -algebra $C(X)$ of all the continuous complex-valued functions on X . The reversal of this process—going from a commutative algebra \mathcal{A} to the topological space whose algebra of functions is \mathcal{A} —is accomplished by the Gel'fand spectrum functor Σ . (Similar to the Stone dualities above, discussed above, Gel'fand duality arises from Hom functors to a dualizing object: \mathbb{C} . $\mathrm{Hom}_{\mathbf{uCC}^*}(-, \mathbb{C})$ is topologized by pointwise convergence; $\mathrm{Hom}_{\mathbf{KHaus}}(-, \mathbb{C})$ is given the uniform norm.) Elements of the C^* -algebra \mathcal{A} can be thought of as continuous complex-valued functions on the space $\Sigma(\mathcal{A})$; the self-adjoint elements are the real-valued ones.

Gel'fand duality has a clear interpretation as a state-observable duality. The objects of the geometric category can be seen as state spaces of classical sorts of systems. Observables, in this analogy, are the continuous real-valued functions on the state space, i.e. the self-adjoint elements of the algebra of observables. The Gel'fand spectrum functor recovers the pure state space from the algebra of observables. We attribute a classical nature to these models since states are associated with well-defined values for all observables simultaneously.

In all these instances, our algebraic categories consist of objects with commutative operations. In quantum theory, the model of a system is specified by a noncommutative C^* -algebra of observables. Understanding the geometric duals of these objects is essential to completing our understanding of how quantum mechanics revises the nature of classical theories and, in particular, notions of states of systems. It is also a fundamental question of purely mathematical interest.

2.3 Quantum State Space

Quantum mechanics challenges the classical notion that all the measurable properties of a physical system have definite values simultaneously. As a most basic example, consider a classical particle. At a given instant in time, it possesses both a position and a velocity. Indeed, once these quantities are given along with a potential energy function associated with the system, classical mechanics yields deterministic predictions about the future of the system. In the most modern and elegant formulations of classical theory, the collection of all obtainable pairings of position quantities and velocity quantities are the points of a geometric object: the state space. The potential energy function provides geometric structure on this collection of points and the time evolution of the system can be described quite simply in terms of special paths within this geometric shape. Measurable quantities are represented as maps which take a point in the shape to their value, i.e. real-valued functions on the space.

The quantum mechanical description of the particle does not ascribe to it both a precise position and a precise velocity. In the standard mathematical formalism of the theory, wherein measurable quantities are represented as self-adjoint operators on a Hilbert space, this is captured by the fact that the operators associated to position and momentum do not commute [41]. This fact that the C^* -algebras of observables associated to quantum systems are noncommutative makes it impossible to frame quantum theory as a geometric theory akin to the above description of classical mechanics. That is, one in which the possible physical states of a system live together in a space and such that the observables are represented by functions which assign quantities to states. This is, essentially, the content of the Kochen-Specker theorem [46], which we review in the next chapter.

However, by admitting more liberal attitudes towards what one considers a ge-

ometric space, one has hope of constructing something which may be justifiably called *quantum state space* in this sense. Of course, this term can be applied to, say, a Hilbert space or to spaces of density matrices. What we seek here is a notion of state such that the algebra of observables is realized as “functions to outcomes” on the state space as opposed to the more general notion outlined in Section 2.1. Phrased alternatively, we seek to construct the geometric objects dual to C^* -algebras; to generalize the Gel’fand spectrum functor from acting on only commutative C^* -algebras to acting on all C^* -algebras.

We have two hints for how to explicitly realize these geometric objects. The first comes from physicists studying the notion of contextuality. Contextuality is a novel phenomenon of quantum theory where the outcome of an experiment measuring an observable depends on the procedure used for measuring that observable. Geometries invented to study this phenomenon give us a starting point. The second hint comes from pure mathematicians who have been studying the geometrical dual of C^* -algebras for many decades. Synthesizing the insights of these bodies of scholarship sheds light on each and constitutes the overall contribution of the work described in this dissertation.

2.4 The Noncommutative Geometry of C^* -algebras

Noncommutative geometry is the mathematical study of noncommutative algebras by the extension of geometric tools which have been rephrased in the language of commutative algebra [45]. Given a duality between geometric objects and commutative algebras, like Gel'fand duality, we can rephrase geometric concepts by expressing them algebraically in terms of functions. For example, if we wish to algebraically express the idea of an open set of a topological space X , we might think about the set of functions which vanish outside of it and note that this set is an ideal in $C(X)$. In fact, there is a bijective correspondence between closed ideals of $C(X)$ and open sets of X . As a more complicated example, the Serre-Swan theorem [67] allows us to identify vector bundles over X with finitely generated projective $C(X)$ -modules. Remarkably, these algebraic descriptions of geometric concepts do not rely crucially on the commutativity of $C(X)$. This allows us to generalize geometric tools and intuition to noncommutative algebras \mathcal{A} by using these same algebraic descriptions. This justifies thinking of a noncommutative C^* -algebra as a *noncommutative topological space*. The elements of the C^* -algebra \mathcal{A} can be thought of as continuous complex-valued functions on a metaphorical noncommutative space. Such a space defies explicit description by conventional mathematical ideas about what a space is; for example, it cannot be thought of as a collection of points for such an object always has a commutative algebra of functions.

One of the best examples of an extension of a topological tool to the setting of noncommutative spaces is that of K -theory. The isomorphism classes of vector bundles over a space X form a semigroup under direct sum and the Grothendieck group of this semigroup is $K(X)$. The K functor is an important cohomological invariant in the study of topology. By using the geometry-to-algebra dictionary described above, we can define an extension of K to C^* -algebras \mathcal{A} in terms of equivalence classes of finitely generated projective \mathcal{A} -modules which is called K_0 .

It is an extension in the sense that when \mathcal{A} is commutative, i.e. $\mathcal{A} \simeq C(X)$ for a space X , then $K_0(\mathcal{A}) \simeq K(X)$. In this way, we obtain a most powerful invariant of C^* -algebras [24]; one which is the basis of a classification program. We note that in the modern account of operator K_0 , one uses an equivalent formulation in terms of equivalence classes of projections in matrix algebras over \mathcal{A} [60].

With considerable effort, this process of translation from geometry to algebra yields a conceptual dictionary covering a vast terrain within mathematics. It is not just topological concepts which can be translated into the language of algebra; there exist noncommutative extensions of measure theory, differential geometry, etc. [15]

Geometry

Algebra


continuous function from a space to \mathbb{C}	element of the algebra (operator)
continuous function from a space to \mathbb{R}	self-adjoint element of the algebra
range of a function	spectrum of an operator
open set	closed, 2-sided ideal
vector bundle	finite, projective module
cartesian product	minimal tensor product
disjoint union	direct sum
infinitesimal	compact operator
Borel probability measure	state
integral	trace
1-point compactification	unitalization
...	...

2.5 The ‘Geometry’ in Noncommutative Geometry

The unreasonable effectiveness of topological tools and intuition in the study of C^* -algebras suggests the existence of a deeper principle at work. The method of translating geometric ideas into algebra in order to generalize them is powerful but strikes one as somewhat difficult and clumsy. Ideally, one would hope for a new conception of space, of which the commutative/topological situation is a special case, which serves as the categorical dual of noncommutative C^* -algebras. That is, such a space would extend the notion of the Gel’fand spectrum of a commutative algebra to the noncommutative case and assign to an algebra \mathcal{A} an object whose set of continuous functions is, in some sense, \mathcal{A} . As pointed out above, an explicit description of (currently imaginary) noncommutative topological spaces is very difficult since such spaces defy most contemporary ideas about mathematical spaces. It is difficult to know how to begin defining such an object. However, we can imagine that equipped with such an explicit description, should it not depart too far from the commutative situation, we could find far more natural and intuitive methods of extending topological tools.

We draw inspiration from the Isham-Butterfield approach to contextuality [34]. In this approach, a quantum system with separable, unital von Neumann algebra of observables \mathcal{A} is analyzed in terms of a presheaf of the Gel’fand spectra of the unital, commutative sub-von-Neumann-algebras of \mathcal{A} . The commutative subalgebras are connected by inclusion maps and so their spectra are connected by restriction maps induced by Σ . The elements of these spectra play the role of valuations of a collection of jointly measurable observables, i.e. contextual outcomes. As we show in Chapter 3, canonical quantum states are easily recovered as the limit of this presheaf after composing with the probability distribution functor. In Chapter 4 and Chapter 5, we modify this presheaf in order to construct an automatic method of translating topological concepts from geometry to algebra.

The criterion for a successful spatial manifestation of noncommutative space is that it naturally leads to extensions of topological concepts which agree with well-known noncommutative geometric invariants. In effect, we aim to complete the following conceptual commutative diagram:



This informal diagram requires some explanation. The top row describes the two dually equivalent mathematical formalisms for encapsulating the operational content of a classical system: the topological picture, in which states are taken as the primitive concept, and the commutative C^* -algebraic picture, in which observables are taken as primitive.

The arrows give methods for the translation and generalization of concepts. The Gel'fand spectrum functor allows for any notion or theorem phrased in terms of the topological structure of spaces to be translated into algebraic terms; i.e. open sets of a space becomes closed, two-sided ideals of an algebra. Once a concept has been phrased in terms of algebra, it can be applied without modification to the noncommutative case; i.e. finitely-generated projective modules of a commutative algebra (the equivalent of vector bundles) becomes finitely-generated projective modules of a not-necessarily commutative algebra. Thus, the composition of the top and right arrows can be seen as the process of generating the basic entries of the noncommutative dictionary.

A topological concept can be translated in several different ways which means that intuition and judgement must be deployed when determining appropriate algebraic analogues. As a very trivial example, open sets of a space X are in correspondence with both the closed, left ideals of $C(X)$ and the closed, two-sided ideals of $C(X)$ as these two collections are identical in the commutative case. Thus, finding a completely automatic method of translation which eliminates such ambiguities would in itself constitute an advance in the structural understanding of noncommutative geometry.

Akemann and Pedersen [56] proposed to replace the translation process by working directly with Giles-Kummer's [30] and Akemann's [3] noncommutative generalizations of the basic topological notions of open and closed sets. In contrast, we do not employ algebraic generalizations of basic topological notions, but rather, we work with objects which slightly generalize the notion of topological space and come readily equipped with an alternative to the translation process.

In addition to the work of Akemann-Pedersen and Giles-Kummer on noncommutative generalizations of Gel'fand duality, there have been a number of alternative approaches by authors including Alfsen [4], Bichteler et al. [11], Dauns-Hofmann [16], Fell [25], Heunen et al. [40], Kadison [43], Kruml et al. [47], Kruszyński-Woronowicz [48], Mulvey [55], Resende [57], Schultz [63], and Takesaki [68]. An excellent discussion of many of these works is contained in a paper by Fujimoto [28].

Our goal with this work is to solve the above diagram for the mathematical structure \mathbf{G} . The first motivation is to give a geometric manifestation for a notion of noncommutative space (the quantum state space described above) whose existence is currently understood as being merely metaphorical. The second is to exploit this geometric manifestation to give a canonical method for importing concepts of topology to noncommutative algebra.

The primary desideratum of a guess for \mathbf{G} is that it comes equipped with natural methods of generalizing notions from topology and translating them to noncommutative algebra; that is, labels for left and bottom arrows. That the composition of these two arrows match the noncommutative dictionary is what would justify thinking of \mathbf{G} as the geometric manifestation of a noncommutative algebra. Our guess for \mathbf{G} , as inspired by Isham and Butterfield's work, is to consider diagrams of topological spaces. The framework of extensions, developed in Section 4, formalizes how various ways of associating diagrams of topological spaces to noncommutative algebras come with such left and bottom arrows and in this way, yield a noncommutative counterpart for every topological concept. In Chapter 6, we solve for the appropriate \mathbf{G} such that the associated extension of topological K -theory essentially matches up with the established noncommutative K -theory. In Chapter 7, as a verification of this construction of \mathbf{G} , we also use it to extend the notion of open set to the notion of closed, 2-sided ideal.

Chapter 3

Contextuality

The Kochen-Specker theorem and the phenomenon of contextuality in quantum mechanics is described. The spectral presheaf construction of a context-indexed family of state spaces is given an elementary presentation with emphasis on its role as a generalized state space. It is shown how its lack of points is equivalent to the Kochen-Specker theorem while Gleason's theorem is expressed as a correspondence of probability distributions on the spectral presheaf with quantum states.

3.1 Ontology and Epistemology

The famed (Bell-)Kochen-Specker theorem [46] is a result of profound metaphysical significance; one best viewed not as a result strictly concerning quantum theory itself but rather as one concerning all physical theories which aim to provide a more refined explanation of empirically observed phenomena than quantum theory. It asserts the impossibility of constructing a physical theory which reproduces the highly-verified experimental predictions of quantum mechanics while maintaining a commonly-held conception of *realism*.

Intuitively, realism is the notion that physical systems have objective properties which possess an existence independent of any observer or observation; we will formalize this idea shortly. The theorem asserts that the concession which the realistically-minded physicist must make is accepting the phenomenon of contextuality; that is, allowing that the outcome of an experiment may depend not simply on the system's state and the quantity being measured but also on the choice of quantities which are simultaneously measured, or, equivalently, on the choice on experimental procedure used to ascertain the quantity in question. The necessity of this concession is in tension with the assumption that the process of measurement is the benign extraction of a pre-existing property of a system. Herein lies a clear and unavoidable departure of quantum theory from the classical worldview. (One interpretation of the Kochen-Specker theorem available to those who wish to hold onto a classical form of reality underlying quantum theory is to reject only the absolute reductionism which denies the possibility of interactions between the observed system and measuring apparatus affecting the outcome of observation [61].)

An operational theory collating empirical predictions is an inherently epistemological description of a physical system. There is no implication that the states in such a description correspond to all the possible real physical configurations or

that observables correspond to real physical properties. Only a collection of probability distributions which describe our knowledge about the outcomes of possible experiments is provided. Given such a collection of probability distributions, one is immediately drawn to inquiring as to their origin. Is there a realistic underlying picture, as is implicit in classical theory, of the system existing in one of potentially many states, knowledge of which would specify well-defined properties which are independent of observation? Can any indeterminism in the theory be attributed to ignorance of this precise state as in statistical mechanics? The natural framework for addressing these questions are provided by *hidden variable theories* or *ontological models*. [10, 64]

3.1.1 Ontological models

Definition 3.1.1. An ontological model for an operational theory (P, M, F) is a measurable space (Λ, Σ) together with:

1. a map $P \rightarrow D(\Lambda) :: p \mapsto \mu_p$ of preparations to probability distributions on the set of ontic states Λ
2. a map of measurements $m \mapsto (r_m : \Lambda \rightarrow D(O))$ to response functions r_m which yield, for each ontic state λ , a probability distribution $r_m(\lambda)$ on the space of outcomes

such that for any measurement m , preparation p , and event e

$$f(e|m, p) = \int r_m(\lambda)(e) d\mu_p(\lambda)$$

Note here that the we have allowed for the possibility that the ontic state does not fully determine the outcome of every measurement. There are two possible interpretations of this. The first would be that the description of the system in

question encapsulated by the ontic state is simply incomplete. Another would be to assume that the measuring equipment is noisy.

Every operational theory (P, M, F) has a trivial operational ontological model where the ontic state space is simply taken to be the set of preparations, i.e. $(\Lambda, \Sigma) = (P, \mathcal{P}(P))$. The probability distribution associated to a preparation $p \in P$ is simply the Dirac delta distribution for p whereas the the response functions for measurements $m \in M$ are given by $r_m(p)(e) = f(e|p, m)$. So, the mere existence of an ontological model for an operational theory is meaningless; what is interesting is knowing the existence of models which exhibit certain properties.

Insisting that the ontic states determine with certainty the resulting observation of any measurement procedure performed on the system is mathematically formalized by insisting that the distributions $r_m(\lambda, -)$ are Dirac delta distributions. That is, the response of a measurement on a system in a known ontic state is determined.

Definition 3.1.2. A deterministic ontological model for an operational theory (P, M, F) is a measurable space (Λ, Σ) together with:

1. a map $P \rightarrow D(\Lambda) :: p \mapsto \mu_p$ of preparations to probability distributions on the set of ontic states Λ
2. a map of measurements $m \mapsto (r_m : \Lambda \rightarrow O)$ to deterministic response functions

such that for any measurement m , preparation p , and event e

$$f(e|m, p) = \mu_p(r_m^{-1}(e))$$

We saw above that every operational theory admits an ontological model. In fact, if we assume that the space of outcomes is a standard probability space (an assumption covering any reasonable physical situation) and restrict ourselves to modelling, at most, countably infinitely many preparation procedures, we can always find a deterministic ontological model [49].

Theorem 3.1.3. *Every operational theory (P, M, F) where the set of preparations P is countable admits a deterministic model.*

Proof. Take the space of ontic states Λ to be a copy of the unit interval for each preparation: $P \times [0, 1]$. For every preparation $p \in P$, there is an isomorphism (modulo null sets) of probability spaces $\Phi_p : ([0, 1], l) \rightarrow (O, f(-|p, m))$ from the unit interval with Lebesgue measure to the outcome space with the probability distribution associated to measuring m on a preparation p given by the physical theory. The deterministic response functions are then given by $r_m(p, x) = \Phi_p(x)$. \square

3.2 Contextuality in quantum theory: the Kochen-Specker Theorem

The famous no-go result of Kochen and Specker [46], anticipated by Bell [10], asserts the impossibility of embedding the (non-Boolean) lattice of propositions about a quantum system into a Boolean algebra. Phrased in the modern and, perhaps, more intuitively clear framework of ontological models, it is equivalent to asserting the impossibility of constructing a certain type of ontological model for quantum theory.

Suppose we have a quantum system whose algebra of observables is represented by $\mathcal{B}(\mathcal{H})$ where $\dim \mathcal{H} > 2$ (or, indeed, represented by any noncommutative, separable von Neumann algebra without summands of type I_1 or I_2 factor [18]). The measurements M are given by the self-adjoint operators in the algebra of observables. The preparations P are a nonempty collection of pure states. The distributions on outcomes F are given by the Born rule.

Definition 3.2.1. *An operational quantum theory is an operational theory (P, M, F) arising as above.*

That is, an operational quantum theory is simply the empirical predictions of quantum theory. We arrive at our mathematical formalization of the concept of noncontextuality which, in this case, is a property of ontological models for operational quantum theories.

Definition 3.2.2. *A ontological model for operational quantum theory is noncontextual if, whenever $A = f(B)$ for self-adjoint operators $A, B \in M$ and a measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$, then the response functions for A and B are related by f :*

$$r_A = f_* \circ r_B$$

Here, the expression $A = f(B)$ is meant in the sense of the Borel functional calculus and the function f_* acts on distributions on the real line by pushing it forward by f . When the ontological model in question is deterministic, this expression reduces to simply: $r_A = f \circ r_B$. (Traditional formulations of noncontextuality in quantum theory usually begin by assuming a deterministic model. A rigorous derivation of the functional composition principle for deterministic valuations in quantum theory from the basic premise of noncontextuality is given in [36].)

Theorem 3.2.3 (Kochen-Specker, 1967). *No operational quantum theory admits a noncontextual, deterministic ontological model.*

Proof. We briefly summarize the proof for the case where the system in question is described by $\mathcal{B}(\mathcal{H})$ for \mathcal{H} of dimension 3. The case of higher dimensions follows from this and the more general von Neumann algebraic case is treated in [18]. Suppose that Λ is a noncontextual, deterministic model for an operational quantum theory (P, M, F) . Fixing a $\lambda \in \Lambda$, we define a function $\lambda(p)$ on projections $p \in \mathcal{B}(\mathcal{H})$ by $\lambda(p) = r_p(\lambda)$. Since projections are fixed by being squared, by noncontextuality, we find that $\lambda^2 = \lambda$ and thus that the function λ can take only the values 0 or 1. We can assume that λ assigns 1 to the identity operator as this must hold for almost all λ (with respect to a distribution coming from any preparation) for the model to reproduce the Born rule. Since projections are mapped to their orthogonal complements by the function $x \rightarrow 1 - x$, we see, again, by noncontextuality, that λ assigns 1 to precisely one of p and p^\perp . A simple induction argument demonstrates that λ is thus additive on sets of orthogonal projections. Kochen and Specker prove that such a valuation on projections is impossible to construct by providing a collection K of 117 vectors in \mathcal{H} such that there is no subset of K intersecting every orthogonal triple in K precisely once. \square

3.3 The spectral presheaf

The geometric reformulation of the Kochen-Specker theorem in terms of the lack of global sections of a presheaf of topological spaces, introduced by Hamilton, Butterfield, and Isham [34], is the inspiration for the choice of ansatz described in Section 2.5 which guides the work described in the sequel.

The Kochen-Specker theorem precludes a model of the theoretically predicted and empirically observed statistics of quantum theory in terms of a (deterministic) classical state space if one insists that the mathematical representation of an observable is independent of the method used to measure it. The spectral presheaf can be interpreted as a collation of classical sample spaces which are associated to different measurement procedures.

A single experimental procedure may yield outcomes for multiple observables. As a simple consequence of the operational definition of observables, experimental procedures are operationally characterized by the maximal collection of observables for which they yield outcomes. The mathematical formalism of quantum theory encodes the joint measurability of a collection of observables algebraically: a collection of observables which can be measured simultaneously is simply a set $C \subset \mathcal{A}$ of self-adjoint operators which is pairwise commutative. When performing a procedure which yields outcomes for every observable in C , simple classical post-processing yields outcomes for observables which can be expressed as polynomials in C and the limits of such polynomials. This justifies formally representing operational equivalence classes of experimental procedures in quantum theory as closed, commutative $*$ -subalgebras of an algebra of observables.

Definition 3.3.1. *A context of a unital von Neumann algebra \mathcal{A} is a unital, commutative sub-von-Neumann-algebra of \mathcal{A} . The context category $\mathcal{C}(\mathcal{A})$ is the subcategory of commutative, unital von Neumann algebras whose objects are the contexts of \mathcal{A} and*

whose morphisms are the inclusion maps between them.

For every context C , the Gel'fand spectrum functor can be used to construct a sample space $\Sigma(C)$ whose points represent the possible outcomes for a procedure jointly measuring all the observables in C . The elements o of $\Sigma(C)$ are functions which assign real numbers to the observables in C and which preserve addition and multiplication. (These conditions are easily justified on physical grounds and are sufficient to guarantee that o assigns to a self-adjoint operator Q an element of the spectrum of Q .) This collection of functions comes equipped with an extremally disconnected topology coming from pointwise convergence which is discrete in the case that \mathcal{A} is finite-dimensional.

Definition 3.3.2 (Spectral presheaf). *Suppose \mathcal{A} is a unital, separable von Neumann algebra without type I_1 or I_2 summands. The spectral presheaf is the contravariant functor from $\mathcal{C}(\mathcal{A})$ to \mathbf{KHaus} which maps each object and morphism of $\mathcal{C}(\mathcal{A})$ to its image under the Gel'fand spectrum functor.*

An inclusion map $\iota : C_1 \hookrightarrow C_2$ goes from a context C_1 which represents a procedure measuring a subset of the observables measured by the procedure represented by C_2 , i.e. a course-graining. The image under the Gel'fand spectrum functor of such an inclusion $\Sigma(\iota) : \Sigma(C_2) \rightarrow \Sigma(C_1)$ acts by restriction: an outcome $o \in \Sigma(C_2)$ is mapped to $o|_{C_1}$.

Thus, a global section of the spectral presheaf of \mathcal{A} is a choice of $o_C \in \Sigma(C)$ for all contexts C of \mathcal{A} such that $o_{C_1} = o_{C_2}|_{C_1}$ whenever $C_1 \subset C_2$.

3.3.1 Points: the Kochen-Specker theorem

Theorem 3.3.3. *The Kochen-Specker theorem is logically equivalent to the non-existence of global sections of spectral presheaves.*

Proof. Suppose the Kochen-Specker theorem is false and Λ is a noncontextual, deterministic ontological model for the operational quantum theory (P, M, F) of \mathcal{A} . Fixing a $\lambda \in \Lambda$, denote by $o(m)$ the value of the response function $r_m(\lambda)$ associated to a self-adjoint operator $m \in M$ which, by determinism, is a definite real number. For each context $C \subset \mathcal{A}$, define $o_C : C \rightarrow \mathbb{R}$ by restricting o to the self-adjoints in C and extending by linearity to all of C . That $o_{C_1} = o_{C_2}|_{C_1}$ whenever $C_1 \subset C_2$ is trivially satisfied; it remains to show that $o_C \in \Sigma(C)$, i.e., they are linear and multiplicative functions. Let A and B be self-adjoints in C . Since they commute, there is a self-adjoint element $S \in C$ such that $A = f(S)$ and $B = g(S)$. Thus, if $h = f + g$, then $A + B = h(S)$. Therefore,

$$\begin{aligned}
o_C(A + B) &= o_C(h(S)) \\
&= r_{h(S)}(\lambda) \\
&= h \circ r_S(\lambda) \\
&= f \circ r_S(\lambda) + g \circ r_S(\lambda) \\
&= r_{f(S)}(\lambda) + r_{g(S)}(\lambda) \\
&= r_A(\lambda) + r_B(\lambda) \\
&= o_C(A) + o_C(B)
\end{aligned}$$

A similar proof establishes multiplicativity. Therefore, the functions o_C constitute a global section of the spectral presheaf associated to \mathcal{A} .

Now, suppose that o_C are the sections of a global section of the spectral presheaf of \mathcal{A} . Denote by o the valuation on all projections of \mathcal{A} given by $o(p) = o_C(p)$ for any context C containing p (which is well-defined by the restriction condition since $\mathbb{C}p + \mathbb{C}p^\perp \subset C$). It assigns to each p one of its eigenvalues: either 0 or 1. It is finitely additive on any set S of orthogonal projections since the set S generates a context C_S and o_{C_S} is additive. Therefore, by Gleason's theorem [31], there is a state ρ which extends o . The state ρ must be pure as a proper convex combination of states could not yield deterministic valuations on all projections. Thus, one can construct

a trivial ontological model which is noncontextual and deterministic with a single ontic state for the system represented by \mathcal{A} and the single preparation ρ . \square

Thus, the impossibility of providing a mathematical model in the classical sense for quantum theory is expressed by constructing a geometric object associated to a quantum system by collating the sample spaces associated to contexts, linked by a simple consistency condition related to course-graining, and demonstrating that said object possesses no ‘global points’.

3.3.2 Distributions: Gleason’s theorem

These geometries represented by spectra presheaves, do, however, possess global probability distributions. Remarkably, these distributions are in correspondence with (possibly mixed) quantum states. Just as the lack of points of spectral presheaves is equivalent to a landmark theorem of quantum foundations (the Kochen-Specker theorem), the correspondence of distributions on spectral presheaves with quantum states is equivalent to Gleason’s theorem [31, 33]. This observation is easily made and succinctly expressed using the framework described in following chapters and was originally made by de Groote [17]; here, we give an elementary description to emphasize its physical interpretation.

Suppose \mathcal{A} is an algebra of observables for a quantum system as above and ρ is a mixed state of this system (i.e, a positive normalized linear functional). For every context $C \subset \mathcal{C}$, ρ supplies a Borel, regular probability measure on the space of outcomes $\Sigma(C)$. This is simply the content of the Riesz-Markov-Kakutani representation theorem [44, p53] as applied to the functional $\rho|_C$. (We will use $\rho|_C$ to also denote the corresponding measure on $\Sigma(C)$.) This distribution is simply that one which is predicted by quantum theory for measurements of the observables in C , measured by the procedure represented by C , of a system in the state ρ . When does

a context-indexed family of probability distributions P_C on $\Sigma(C)$ arise in such a way from a mixed state?

When a context C_1 contains fewer observables than a context C_2 , i.e. $C_1 \subset C_2$, then we can derive a distribution P_1 on $\Sigma(C_1)$ from a distribution P_2 on $\Sigma(C_2)$ in a particularly simple way: marginalization. The likelihood of an event $e \subset \Sigma(C_1)$ should be the likelihood of the event in $\Sigma(C_2)$ consisting of those outcomes which assign to the observables in C_1 all the same values as some outcome in e . This is expressed equationally as:

$$P_1(e) = P_2(R^{-1}(e))$$

Here, R is $\Sigma(\iota) : \Sigma(C_2) \rightarrow \Sigma(C_1)$ which the map described above which restricts a real-valued function of C_2 to a function of C_1 .

Definition 3.3.4. *A context-indexed family of distributions P_C on $\Sigma(C)$ is consistent when the above marginalization condition holds for every pair of contexts $C_1 \subset C_2$.*

Every family of distributions provided by a quantum state ρ is consistent. To prove this, consider an elementary quantum event represented by a projection P in C_1 and denote by $E_P^{C_1} \subset \Sigma(C_1)$ the set of outcomes where a measurement of P in the context C_1 returns one, i.e. those functions assigning to P the value one. An outcome $\Sigma(C_2)$ assigns to P the value one if and only if its restriction to C_1 does and so $E_P^{C_2}$ is precisely $R^{-1}(E_P^{C_1})$. The measurement of P in either context returns one with a likelihood of $\rho(P)$ and, therefore,

$$\begin{aligned} \rho|_{C_1}(E_P^{C_1}) &= \rho(P) \\ &= \rho|_{C_2}(E_P^{C_2}) \\ &= \rho|_{C_2}(R^{-1}(E_P^{C_1})). \end{aligned}$$

As this equation holds for all the elementary quantum events of C_1 , which form a basis for the topology of $\Sigma(C_1)$, it holds for all events. Thus, the context-indexed

family of distributions $\{\rho|_C\}$ is consistent. (A similar proof of this fact appears in [1]).

We demonstrated above that any quantum state yields a consistent family; now we show that from a consistent family, we can construct the unique state which gives rise to it. The proof relies on Gleason's theorem [31, 33] and, indeed, provides an alternative interpretation of that landmark result.

Suppose we are given a family of probability distributions $\rho|_C$ on $\Sigma(C)$ for every context $C \in \mathcal{A}$ which is consistent. If P is a projection, C is any context which contains P , and $E_P^C \subset \Sigma(C)$ is the event consisting of outcomes which assign to P the value one, then $\rho|_C(E_P^C) = \rho|_{C_P}(E_P^{C_P})$ where C_P is the context generated by P . This follows from consistency and the fact that $C_P \subset C$. We can thus define μ by $\mu(P) \equiv \rho|_C(E_P^C)$.

That μ assigns one to the identity projection is immediate. Now, if P_1 and P_2 are orthogonal, denote their sum by P_3 and the context generated by P_1 and P_2 by C' . Since, by linearity, an outcome $o \in \Sigma(C')$ assigns one to P_3 if and only if the sum of the values o assigns to P_1 and P_2 is one, it follows that the event $E_{P_3}^{C'}$ is the disjoint union of $E_{P_1}^{C'}$ and $E_{P_2}^{C'}$. By additivity of the measure $\rho|_{C'}$:

$$\begin{aligned} \mu(P_3) &= \rho|_{C'}(E_{P_3}^{C'}) \\ &= \rho|_{C'}(E_{P_1}^{C'}) + \rho|_{C'}(E_{P_2}^{C'}) \\ &= \mu(P_1) + \mu(P_2) \end{aligned}$$

So, by Gleason's theorem, we can extend μ to a state M whose derived family of distributions coincide with the family $\{\rho|_C\}$ on elementary—and thus on all—events.

Chapter 4

Spatial Diagrams

We introduce the technical machinery necessary for contravariantly functorially associating diagrams of topological spaces, representing quotient spaces of a noncommutative space, to noncommutative C^* -algebras. We axiomatize those functors which associate to a C^* -algebra a diagram whose objects are spectra of contexts and whose morphisms are such that the association yields a natural method of extending functors (described in the next chapter) which act on compact Hausdorff spaces to functors which act on all unital C^* -algebras. The specific choice of morphisms leading to the correct method of extension, as determined by the results of Chapters 6 and 7, is given by those which arise from restricting inner automorphisms. A simple example is presented.

4.1 The Categories of all Diagrams in \mathcal{C}

We propose to associate to each unital C^* -algebra \mathcal{A} a diagram of topological spaces whose objects are the spectra of the unital, commutative sub- C^* -algebras of \mathcal{A} . Such an association should be expected to be (contravariantly) functorial. Typically, one thinks of a diagram $D : J \rightarrow C$ in a category C as living inside the functor category C^J . This is inadequate for our purposes as different algebras will have different sets of commutative sub- C^* -algebras and thus be associated to diagrams of different shapes. We introduce a very general construction which allows considering diagrams of different shapes on the same footing.

Definition 4.1.1. *For any category C , $\underline{Diag}(C)$, the covariant category of all diagrams in C , has as objects all the functors D from any small category S to C . Morphisms from $D_1 : S_1 \rightarrow C$ to $D_2 : S_2 \rightarrow C$ are given by pairs (f, η) where f is a functor from $S_1 \rightarrow S_2$ and η is a natural transformation from D_1 to $D_2 \circ f$.*

The contravariant category of all diagrams $\overleftarrow{Diag}(C)$ has all contravariant functors to C as objects; the morphisms from D_1 to D_2 are pairs (f, η) where f is a functor from $S_2 \rightarrow S_1$ and η is a natural transformation from $D_1 \circ f$ to D_2 .

(These can be constructed by considering the colax-slice and lax-slice 2-categories Cat / C [62] and forgetting the 2-categorical structure.) The composition $(g, \mu) \circ (f, \eta)$ of two $\underline{Diag}(C)$ -morphisms is given by $(gf, (\mu f)\eta)$ where $(\mu f)_a$ is $\mu_f(a)$ and the composition of natural transformations is componentwise. Note that if F is a functor from C to C' , F naturally induces a functor from $\underline{Diag}(C)$ to $\underline{Diag}(C')$ which we will also denote by F . Explicitly, if $D : A \rightarrow C$, then $F(D)$ is simply $F \circ D$. For a $\underline{Diag}(C)$ -morphism (f, η) , F sends (f, η) to the $\underline{Diag}(C')$ -morphism $(f, F\eta)$ where $(F\eta)_a$ is $F(\eta_a)$. The functor F also induces, in a similar fashion, a functor from $\overleftarrow{Diag}(C)$ to $\overleftarrow{Diag}(C')$. If F is contravariant, then it induces a contravariant functor from $\underline{Diag}(C)$ to $\overleftarrow{Diag}(C')$ and one from $\overleftarrow{Diag}(C)$ to $\underline{Diag}(C')$.

4.2 Semispectral Functors

Having defined a category which can simultaneously accommodate diagrams of varying shapes, we are ready to begin defining our contravariantly functorial associations of diagrams of topological spaces to C^* -algebras. We will define a class of such functorial associations. What all these functors from the category of unital C^* -algebras to diagrams of compact Hausdorff spaces have in common is that they will associate to each unital C^* -algebra a diagram (i.e. a functor) with domain a subcategory of the category of unital, commutative C^* -algebras. In fact, in each case, the objects of the domain subcategory of the diagram associated to a C^* -algebra are its unital, commutative sub- C^* -algebras. The class of morphisms in the domain subcategory, however, will be allowed to vary.

Our motivating example is the spectral presheaf. The recipe for its construction which we aim to generalize is:

1. take a unital von Neumann algebra \mathcal{A}
2. consider the subcategory $s(\mathcal{A}) \subset \text{uCVnA}$ of unital, commutative von Neumann algebras with unital $*$ -homomorphisms whose objects are the unital, commutative sub-von-Neumann algebras (contexts) of \mathcal{A} and whose morphisms are the inclusions between such subalgebras
3. consider the inclusion functor $\iota_{\mathcal{A}}$ of this subcategory of unital, commutative von Neumann algebras into unital, commutative C^* -algebras; this is an object of $\underline{\text{Diag}}(\text{uCC}^*)$
4. compose the Gel'fand spectrum functor with this inclusion functor to yield an object of $\underline{\text{Diag}}(\text{KHaus})$.

This association of spectral presheaves to algebras can be made functorial in a natu-

ral way. For a $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$, we can define a $\underline{Diag}(\mathbf{uCC}^*)$ -morphism (f, η) . The functor $f : s(\mathcal{A}) \rightarrow s(\mathcal{B})$ sends a subalgebra $U \subset \mathcal{A}$ to $\phi(U) \subset \mathcal{B}$ and an inclusion $U \hookrightarrow V$ to the inclusion $\phi(U) \hookrightarrow \phi(V)$. The natural transformation $\eta : i_{\mathcal{A}} \rightarrow i_{\mathcal{B}} \circ f$ has components η_U given by $\phi|_U : U \rightarrow \phi(U)$. The Gel'fand spectrum functor Σ yields a functor $\Sigma : \underline{Diag}(\mathbf{uCC}^*) \rightarrow \underline{Diag}(\mathbf{KHaus})$ and the image under this functor of (f, η) is a $\underline{Diag}(\mathbf{KHaus})$ -morphism between the spectral presheaves of \mathcal{A} and \mathcal{B} .

We will generalize this recipe to C^* -algebras. However, we will also want to consider other choices of which morphisms to include in our diagrams. In the next chapter, we see that an association of diagrams of space to algebras automatically yields a method of extending topological functors to functors which act on all unital C^* -algebras. The family of morphisms we include in our diagrams affect the resulting method of extensions. Thus, we vary the family of morphisms in order to solve for the one whose method of extending functors matches up with the canonical generalization process of noncommutative geometry; this was the motivation behind the reconstruction of the definition of operator K -theory.

Definition 4.2.1. A functor $\sigma : \mathbf{uCC}^* \rightarrow \underline{Diag}(\mathbf{uCC}^*)$ is called semispectral if:

(1) $\sigma(\mathcal{A})$ is an inclusion functor from a subcategory $\text{dom}(\sigma(\mathcal{A}))$ of \mathbf{uCC}^* whose objects are the unital, commutative sub- C^* -algebras of \mathcal{A}

(2) For a unital $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$, $\sigma(\phi)$ is (f, η) where f is a functor from $\text{dom}\sigma(\mathcal{A})$ to $\text{dom}\sigma(\mathcal{B})$ which takes a unital, commutative sub- C^* -algebra $U \subset \mathcal{A}$ to $\phi(U)$ and η is the natural transformation which associates to U the $*$ -homomorphism $\phi|_U : U \rightarrow \phi(U)$

(3) If \mathcal{A} is commutative, then it is terminal in $\text{dom}\sigma(\mathcal{A})$

The third condition is essential for constructing an extension process associated to a semispectral functor.

4.3 Spatial diagrams

We will be working primarily with spatial diagrams, which is a particular way of associating diagrams of topological spaces to unital C^* -algebras using a semispectral functor. We will define this semispectral functor in two steps and then define the spatial diagram functor $G : \mathfrak{u}C^* \rightarrow \underline{Diag}(\mathbf{KHaus})$ by composition with the Gel'fand spectrum functor.

Definition 4.3.1. *The unitary subcategory $S(\mathcal{A})$ of $\mathfrak{u}CC^*$ of a unital C^* -algebra \mathcal{A} has as objects all unital, commutative sub- C^* -algebras of \mathcal{A} and, as arrows, all morphisms which arise as restrictions of inner automorphisms of \mathcal{A} . That is, a morphism $m : U \rightarrow V$ in $\mathfrak{u}CC^*$ (between two unital, commutative sub- C^* -algebras $U, V \subset \mathcal{A}$) is contained in $S(\mathcal{A})$ if and only if there is a unitary element $u \in \mathcal{A}$ such that $uUu^* \subset V$ and $m(a) = uau^*$ for all $a \in U$.*

The composition of two such arrows is given by conjugation by the product of the two original unitaries and so $S(\mathcal{A})$ is indeed a subcategory of $\mathfrak{u}CC^*$. These arrows are all of the form $i \circ r$ where i is an inclusion and r is an isomorphism between subalgebras which are related by unitary rotation.

Definition 4.3.2. *The unitary semispectral functor $g : \mathfrak{u}C^* \rightarrow \underline{Diag}(\mathfrak{u}CC^*)$ sends a unital C^* -algebra \mathcal{A} to the inclusion functor $\iota_{\mathcal{A}} : S(\mathcal{A}) \rightarrow \mathfrak{u}CC^*$. The action of g on unital $*$ -morphisms is fixed by Condition (2) in Definition 4.2.1: a unital $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$, $\sigma(\phi)$ is (f, η) where f is a functor from $S(\mathcal{A})$ to $S(\mathcal{B})$ which takes a unital, commutative sub- C^* -algebra $U \subset \mathcal{A}$ to $\phi(U) \subset \mathcal{B}$ and η is the natural transformation which associates to U the $*$ -homomorphism $\phi|_U : U \rightarrow \phi(U)$*

Note that when \mathcal{A} is commutative, the set of arrows are simply the inclusions which is why the Condition (3) of Definition 4.2.1 holds.

Definition 4.3.3. *The spatial diagram functor $G : \mathfrak{u}C^* \rightarrow \underline{Diag}(\mathbf{KHaus})$ is $\Sigma \circ g$.*

The topological spaces in the diagram $G(\mathcal{A})$ should be thought of as being those which arise as quotient spaces of the hypothetical noncommutative space underlying \mathcal{A} . To see this, note that a sub- C^* -algebra of $C(X)$ yields an inclusion $i : C(Y) \rightarrow C(X)$ which corresponds to a continuous surjection $\Sigma i : X \rightarrow Y$; this surjection is a quotient map since both X and Y are compact and Hausdorff [69, p12]. Thus, in accordance with the central tenet of noncommutative geometry, unital sub- C^* -algebras of a unital noncommutative algebra \mathcal{A} are to be understood as having an underlying noncommutative space which is a quotient space of the noncommutative space underlying \mathcal{A} . By considering only the commutative subalgebras, we are restricting our attention to the tractable quotient spaces: those which are genuine topological spaces. The morphisms of the diagram serve to track how these quotient spaces fit together inside the noncommutative space.

We will require, in Chapter 6, a slight modification of the unitary subcategory:

Definition 4.3.4. *The finitary unitary subcategory $S_f(\mathcal{A}) \subset S(\mathcal{A})$ of uCC^* has as objects all unital, finite-dimensional commutative sub- C^* -algebras of \mathcal{A} and, as arrows, all morphisms which arise as restrictions of inner automorphisms of \mathcal{A} . That is, a morphism $m : U \rightarrow V$ in uCC^* (between two unital, finite-dimensional, commutative sub- C^* -algebras $U, V \subset \mathcal{A}$) is contained in $S(\mathcal{A})$ if and only if there is a unitary element $u \in \mathcal{A}$ such that $uUu^* \subset V$ and $m(a) = uau^*$ for all $a \in U$.*


This is used to define a finitary version g_f of the unitary semispectral functor g :

Definition 4.3.5. *The functor $g_f : \text{uCC}^* \rightarrow \underline{\text{Diag}}(\text{uCC}^*)$ sends a unital C^* -algebra \mathcal{A} to the inclusion functor $\iota_{\mathcal{A}} : S_f(\mathcal{A}) \rightarrow \text{uCC}^*$. For a unital $*$ -morphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$, we define $g_f(\phi)$ to be (f, η) where f is a functor from $S_f(\mathcal{A})$ to $S_f(\mathcal{B})$ which takes a unital, finite-dimensional, commutative sub- C^* -algebra $U \subset \mathcal{A}$ to $\phi(U) \subset \mathcal{B}$ and η is the natural transformation which associates to U the $*$ -homomorphism $\phi|_U : U \rightarrow \phi(U)$.*

4.4 Example: $M_2(\mathbb{C})$

The entire spatial diagram of an arbitrary C^* -algebra is usually rather difficult to compute explicitly. In this section, we describe the simplest case of the 2×2 matrix algebra $M_2(\mathbb{C})$. The topological spaces which are objects of the spatial diagram $G(M_2(\mathbb{C}))$ are themselves the points of a topological space. (The spaces of objects are computed for higher dimensional matrix algebras by Caspers et al. in [14]; they are Grassmannian manifolds.)

There is a single trivial one-dimensional unital sub- C^* -algebra consisting of the scalar multiples of the identity; its spectrum is a single point. The two-dimensional unital sub- C^* -algebras are all of the form $\mathbb{C}p + \mathbb{C}p^\perp$ for a rank 1 projection p ; the spectrum of these algebras are discrete two point spaces. The morphisms of the diagram include a quotient map from each of these two point spaces to the trivial one point space. As all the two-dimensional sub- C^* -algebras are unitarily equivalent, the diagram contains all possible bijections between pairs of two-point spaces (including permutations).



The one-dimensional projections P of $M_2(\mathbb{C})$ can be identified with their image: a line in \mathbb{C}^2 . They are thus parameterized by the points of the complex projective line \mathbb{CP}^1 : equivalence classes of nonzero pairs of complex numbers (α, β) up to rescaling by nonzero complex factors. Two different projections generate the same sub- C^* -algebra precisely when they are orthogonal complements: (α, β) and

$(-\beta, \alpha)$.

The quotient map identifying these points

$$\pi : \mathbb{CP}^1 \rightarrow \mathbb{RP}^1$$

is the antipodal map in the S^2 model of the complex projective line. We thus view the objects of $G(M_2(\mathcal{A}))$ as a space with two connected components. One consists of a single point representing the trivial unital sub- C^* -algebra. The second is \mathbb{RP}^1 and its points are the two-dimensional sub- C^* -algebras. The spectra of these sub- C^* -algebras can be identified with their fibres $\pi^{-1}([P]) = \{P, P^\perp\} \subset \mathbb{CP}^1$.

Chapter 5

Extensions of Topological Functors

The generalization of limit and colimit functors which act on certain functor categories to ones which act on categories of diagrams is given. This allows defining the extension of a topological functor to a noncommutative algebraic one, given an association of diagrams of spaces to algebras as described in the previous chapter. The extension process is interpreted as decomposing a noncommutative space into tractable quotient spaces, applying a topological functor to each one, and pasting together the result. The formulations of the Kochen-Specker theorem and Gleason's theorem described in Chapter 3 are described in this framework.

5.1 The Generalized Limit and Colimit Functors

A key feature of the construction of $\underline{Diag}(C)$ in the case where C is cocomplete is the existence of a generalized colimit functor $\underline{\lim} : \underline{Diag}(C) \rightarrow C$. It assigns to a functor $F : A \rightarrow C$ the same object of C which is assigned to F by the colimit functor of C^A . If η is a natural transformation between F and $G : A \rightarrow C$ then $\underline{\lim}$ assigns to the $\underline{Diag}(C)$ -morphism between F and G given by (id_A, η) the same C -morphism assigned to η by the colimit functor of C^A . What is novel is the ability to assign C -morphisms between colimits of diagrams of different shapes.

Everything in this section applies equally well (that is, all dual statements hold true) to the case where C is complete, in which case we have a generalized limit functor $\underline{\lim} : \underline{Diag}(C) \rightarrow C$.

Recall that the colimit of a functor F from A to a cocomplete category C can be expressed as a coequalizer of two coproducts [50, p355]:

$$\coprod_{u:i \rightarrow j} F(\text{dom } u) \xrightleftharpoons[\tau]{\theta} \coprod_i F(i)$$

The first coproduct is over all arrows $u : i \rightarrow j$ of A and the second is over all objects of A . We denote the canonical injections for these coproducts by

$$\lambda_u : F(\text{dom } u) \rightarrow \coprod_{u:i \rightarrow j} F(\text{dom } u)$$

and

$$\kappa_i : F(i) \rightarrow \coprod_i F(i) .$$

The morphisms θ and τ can be defined by specifying their compositions with the λ_u :

$$\theta \lambda_u = \kappa_{\text{dom } u}$$

$$\tau \lambda_u = \kappa_{\text{cod } u} F(u)$$

The advantage of this coequalizer presentation of the colimit of is that we may determine a C -morphism between the colimits of two functors F and G by specifying

a natural transformation between their coequalizer diagrams. That is, by giving its components; C-morphisms N and M such that the following diagrams commute:

$$\begin{array}{ccc} \coprod_{u:i \rightarrow j} F(\text{dom} u) & \xrightarrow{\theta} & \coprod_i F(i) \\ \downarrow N & & \downarrow M \\ \coprod_{u':i' \rightarrow j'} G(\text{dom} u') & \xrightarrow{\theta'} & \coprod_{i'} G(i') \end{array}$$

$$\begin{array}{ccc} \coprod_{u:i \rightarrow j} F(\text{dom} u) & \xrightarrow{\tau} & \coprod_i F(i) \\ \downarrow N & & \downarrow M \\ \coprod_{u':i' \rightarrow j'} G(\text{dom} u') & \xrightarrow{\tau'} & \coprod_{i'} G(i') \end{array}$$

We denote the canonical injections into the coproducts for G by λ'_u and κ'_i .

Given a $\underline{Diag}(\mathbf{C})$ -morphism (f, η) between F and G we define N and M by giving their compositions with the canonical injections:

$$N\lambda_u = \lambda'_{f(u)}\eta_{\text{dom} u}$$

$$M\kappa_i = \kappa'_{f(i)}\eta_i$$

It is straightforward to verify that the above diagrams commute, that is, that

$$\theta'N = M\theta$$

and that

$$\tau'N = M\tau,$$

by computing the composition of these maps with the λ_u . The C-morphism assigned by $\underline{\lim}$ to (f, η) is then defined to be that morphism which is induced by the natural transformation (whose components are N and M) between the coequalizer diagrams for the colimits of F and G .

Functoriality of $\underline{\lim}$ is then straightforwardly verified by computing the compositions of the components of the natural transformations induced by (f, η) and (g, μ)

and seeing that the resulting natural transformation is the same as the one induced by $(gf, (\mu f)\eta)$.

The generalized colimit construction is best illustrated by the example of abelian groups. Let $F : A \rightarrow \text{Ab}$ and $G : B \rightarrow \text{Ab}$ be two diagrams in $\underline{\text{Diag}}(\text{Ab})$ and (f, η) be a morphism from F to G .

First, we describe the colimit of F in Ab (and thus its image under $\underline{\text{lim}} : \underline{\text{Diag}}(\text{Ab}) \rightarrow \text{Ab}$). Let d be the direct sum of the groups $F(a)$ over all objects a in A . If g is an element of the group $F(a)$, we use the notation $(g)_a$ to indicate the element of d which is g in the a^{th} component and 0 in all others. The colimit of F is d modulo the subgroup generated by the elements $(g)_{a_1} - (F(u)(g))_{a_2}$ where $g \in F(a_1)$ and $u : a_1 \rightarrow a_2$ is an arrow of A .

To define how the functor $\underline{\text{lim}} : \underline{\text{Diag}}(\text{Ab}) \rightarrow \text{Ab}$ acts on (f, η) , it is enough to say how the group homomorphism $\underline{\text{lim}}((f, \eta))$ acts on elements of the colimit of F of the form $[(g)_a]$. The image of such an element under $\underline{\text{lim}}((f, \eta))$ is $[(\eta_a(g))_{f(a)}]$. This is well defined for if $u : a_1 \rightarrow a_2$ identifies $(g)_{a_1}$ with $(F(u)(g))_{a_2}$, then $G \circ f(u)$ identifies $(\eta_{a_1}(g))_{f(a_1)}$ with $(\eta_{a_2}(F(u)(g)))_{f(a_2)}$.

5.2 Extensions of Functors

For a fixed semispectral functor σ , we define a natural method for extending contravariant or covariant functors $F : \mathbf{KHaus} \rightarrow \mathbf{C}$ when \mathbf{C} is cocomplete or complete, respectively. The idea is to use σ to turn an algebra \mathcal{A} into a diagram of commutative algebras, apply the Gel'fand spectrum functor to this diagram to yield a diagram of topological spaces, apply F to yield a diagram in \mathbf{C} , and finally, apply the extended colimit or limit functor. We will primarily be working with the case that σ is the unitary semispectral functor g as in Definition 4.3.2.

Intuitively, one should think of the extension process as decomposing a noncommutative space into its quotient spaces, retaining those which are genuine topological spaces, applying the topological functor to each one, and pasting together the result.

Definition 5.2.1. *For a semispectral functor σ , a cocomplete category \mathbf{C} , and a contravariant functor $F : \mathbf{KHaus} \rightarrow \mathbf{C}$, the σ -**extension of F** , denoted $\tilde{F} : \mathbf{uC}^* \rightarrow \mathbf{C}$, is given by:*

$$\tilde{F} = \varinjlim \circ F \circ \Sigma \circ \sigma$$

Extensions of covariant topological functors using a (contravariant) semispectral functor are defined in the same way except that we use the limit functor $\varprojlim : \mathbf{Diag}(\mathbf{C}) \rightarrow \mathbf{C}$ in place of a colimit.

It is the third property in the definition of semispectral functor—that \mathcal{A} is commutative implies that \mathcal{A} is terminal in the category picked out by $\sigma(\mathcal{A})$ —which is crucial in ensuring that \tilde{F} does indeed extend F . As a consequence of this condition, the diagram $F(\Sigma \circ \sigma(\mathcal{A}))$ has $F(\Sigma(\mathcal{A}))$ as a terminal object. Since, for any diagram with a terminal object, the canonical injection from the terminal object to the colimit is an isomorphism, we have that $\tilde{F}(\mathcal{A}) \simeq F \circ \Sigma(\mathcal{A})$. The second property ensures that

for a homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ between commutative algebras, $\tilde{F}(\phi)$ completes the commutative square formed by these isomorphisms and $F \circ \Sigma(\phi)$. Thus, these isomorphisms define a natural equivalence between $\tilde{F}|_{\mathbf{uCC}^*}$ and $F \circ \Sigma$. We have thus proved that:

Theorem 5.2.2. *For a semispectral functor σ , a cocomplete category \mathbf{C} , and a contravariant functor $F : \mathbf{KHaus} \rightarrow \mathbf{C}$,*

$$\tilde{F}|_{\mathbf{uCC}^*} \simeq F \circ \Sigma$$

Dually, this equation also holds when \mathbf{C} is a complete category and $F : \mathbf{KHaus} \rightarrow \mathbf{C}$ is a contravariant functor.

5.3 Theorems of quantum foundations

Having established the framework of extensions, we demonstrate how they can be used to succinctly express two fundamental theorems of quantum foundations: the Kochen-Specker theorem and Gleason's theorem. The proofs of these reformulations are omitted here as they are essentially contained in Section 3.3. The spectral presheaf functor $\sigma : \mathbf{uVnA} \rightarrow \underline{\text{Diag}}(\mathbf{KHaus})$ is as described by the recipe in Section 4.2. In this section, $\tilde{F} = \varprojlim \circ F \circ \Sigma \circ \sigma$ denotes the σ -extension of a covariant functor $F : \mathbf{KHaus} \rightarrow \mathbf{C}$ to a complete category \mathbf{C} . We also restrict our extensions to the full subcategory of those unital von Neumann algebras \mathcal{A} which are separable and contain no type I_1 or I_2 summands.

Theorem 5.3.1 (Kochen-Specker). *The extension \tilde{I} of the identity endofunctor $I : \mathbf{KHaus} \rightarrow \mathbf{KHaus}$ yields the empty set on noncommutative algebras \mathcal{A} .*

Let $D : \mathbf{KHaus} \rightarrow \mathbf{Set}$ be the Borel probability distribution functor which assigns to a topological space X the set of all Borel probability distributions on X and to a continuous function f the corresponding pushforward map f_* on measures (defined by $f_*(\mu)(e) = \mu(f^{-1}(e))$).

Theorem 5.3.2 (Gleason). *The extension \tilde{D} of the Borel probability distribution functor is naturally equivalent to the functor mapping von Neumann algebras to their set of states (positive normalized linear functionals).*

Chapter 6

Reconstructing Operator K -theory

Topological K -theory, defined in terms of isomorphism classes of vector bundles, and its C^* -algebraic generalization, the operator K_0 functor, defined in terms of isomorphism classes of finite, projective modules of C^* -algebras, are reviewed. A novel definition of operator K_0 is then compared to the extension of topological K -theory. This result fixes the restrictions of inner automorphisms as the appropriate class of morphisms in the diagrams associated to C^* -algebras.

6.1 Topological K -theory

Topological K -theory, invented by Atiyah-Hirzebruch [7] after Grothendieck [32], is an extraordinary cohomology theory. That is, it is a sequence of contravariant functors from \mathbf{KHaus} to \mathbf{Ab} which satisfies all the Eilenberg-Steenrod axioms [23] but the dimension axiom. After early successes, including the solution to the classical problem of determining how many linearly independent vector fields can be constructed on \mathbf{S}^n [2], the subject bloomed to include algebraic and analytic versions. The core idea is to describe the geometry of a space by algebraic information about the possible vector bundles over it. Here, we briefly review its definition. Its generalization to C^* -algebras, operator K_0 , is a key tool of noncommutative geometry and will be outlined in the next section.

Definition 6.1.1. *For a compact Hausdorff space X , the vector bundle monoid $V(X)$ is the set of isomorphism classes of complex vector bundles over X with the abelian addition operation of fibrewise direct sum. That is, $[E] + [F] = [E \oplus F]$. A continuous function $f : X \rightarrow Y$ yields an abelian monoid morphism $V(f) : V(Y) \rightarrow V(X)$ by the pullback of bundles. That is, if $p : E \rightarrow Y$ is a bundle over Y , the bundle f^*E is a bundle over X given by the projection to X of*

$$\{(x, v) \in X \times E : f(x) = p(v)\}$$

*and $V(f)([E]) = [f^*E]$.*

This defines a contravariant functor V from compact Hausdorff spaces to abelian monoids.

Definition 6.1.2. *For an abelian monoid M , the Grothendieck group of M is the abelian group $G(M)$ defined as $M \times M$ modulo the equivalence relation*

$$(a, b) \sim (c, d) \iff \exists e \in M : a + d + e = b + c + e.$$

For an abelian monoid homomorphism $\phi : M \rightarrow N$, the abelian group homomorphism $G(\phi)$ is defined by

$$G(\phi)([(a, b)]) = [(\phi(a), \phi(b))] .$$

The Grothendieck group functor is the universal construction of an abelian group from an abelian monoid in the sense that any homomorphism out of $G(M)$ to any group factors uniquely through the natural homomorphism $a \mapsto [(a, 0)]$ of M to $G(M)$.

Definition 6.1.3. The topological K -functor is the contravariant functor defined by $G \circ V$.

From the topological K functor, one can easily construct the full sequence of functors K_n for $n \in \mathbb{N}$.

Definition 6.1.4. The suspension functor S maps the category of compact Hausdorff space to itself by sending a space X to the quotient space

$$X \times [0, 1] / \{(x, 0) \sim (x', 0) \text{ and } (x, 1) \sim (x', 1) \text{ for all } x, x' \in X\}$$

and a continuous function $f : X \rightarrow Y$ to the map $[(x, t)] \mapsto [(f(x), t)]$.

Definition 6.1.5. Topological K -theory is the sequence of functors K_n defined by $K_n = K(S^{|n|}(X))$.

Bott periodicity [8] provides natural isomorphisms $K_n \simeq K_{n+2}$. We are left with $K_0 = K$ and $K_1 = K \circ S$. Note that topological K -theory additionally possesses a ring structure which does not survive in the noncommutative case.

6.2 Operator K -theory

Here, we outline the generalization of topological K -theory to operator K -theory by the canonical method of noncommutative geometry. We provide the definition and properties of the operator K_0 functor which we will use in our analysis of the extension of the topological K functor. These are basic facts found in any introduction to the subject, e.g. [60], [71], or [26, p181].

In order to generalize a topological concept to the noncommutative case, one must begin with a characterisation in terms of commutative algebra of the topological concept in question. In the case of K -theory, this requires phrasing the notion of a *complex vector bundle over X* in terms of the algebra $C(X)$ of continuous, complex-valued functions on X . This rephrasing is provided by the Serre-Swan theorem:

Theorem 6.2.1 (Serre-Swan, [1959]). *] The category of complex vector bundles over a compact Hausdorff space X is equivalent to the category of finitely generated, projective modules of $C(X)$ [67].*

Recall that a projective \mathcal{A} -module is the direct summand of a free \mathcal{A} -module. Roughly, the module associated to a vector bundle E over X is the set of continuous global sections of E with the obvious operations. This justifies considering a finitely generated, projective (left) \mathcal{A} -module to represent a complex vector bundle over the noncommutative space underlying the C^* -algebra \mathcal{A} .

The canonical translation process of noncommutative geometry suggests, having now in our possession an algebraic characterisation in terms of $C(X)$ of the topological notion of complex vector bundle, that we use it to define its noncommutative generalization. That is, define the Murray-von Neumann semigroup of a C^* -algebra to be the abelian monoid of its finitely generated, projective modules (up to the appropriate notion of equivalence and with an appropriate addition op-

eration). It turns out to be more convenient to work with an algebraic gadget which is equivalent to finitely-generated, projective \mathcal{A} -modules: projections in a matrix algebra $M_n(\mathcal{A})$ over \mathcal{A} . As such a module μ is free, there is another such module μ^\perp such that $\mu \oplus \mu^\perp \simeq \mathcal{A}^n$. We identify the module μ with the projection $p : \mathcal{A}^n \rightarrow \mu$, or rather, the canonical representation of that projection as an element of the matrix algebra $M_n(\mathcal{A})$.

Equipped with our algebraic characterisation of vector bundles, we are ready to begin defining operator K -theory in a manner directly analogous with the construction of topological K -theory.

Definition 6.2.2 (The Murray-von Neumann semigroup for unital \mathcal{A}). *The abelian monoid $V(\mathcal{A})$ is the set*

$$\bigsqcup_{n \in \mathbf{N}} \{p \in M_n(\mathcal{A}) : p \text{ is a projection}\}$$

modulo the Murray-von Neumann equivalence relation defined by identifying $p \in M_n(\mathcal{A})$ with $q \in M_m(\mathcal{A})$ whenever there is a partial isometry v in the C^ -algebra $M_{m,n}(\mathcal{A})$ of $m \times n$ rectangular matrices such that $p = vv^*$ and $q = v^*v$. It is equipped with the abelian addition operation*

$$[p] + [q] = \left[\begin{pmatrix} p & 0 \\ 0 & q \end{pmatrix} \right].$$

A unital $$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ yields a homomorphism $V(\phi) : \mathcal{A} \rightarrow \mathcal{B}$*

$$[p] \mapsto [\phi(p)]$$

where ϕ acts on elements of $M_n(\mathcal{A})$ by entrywise application.

Definition 6.2.3. *The operator K_0 functor is the covariant functor defined by $G \circ V$.*

As in the topological case, one can easily construct the full sequence of functors K_n for $n \in \mathbf{N}$.

Definition 6.2.4. The suspension functor S maps the category of C^* -algebras to itself by sending a C^* -algebra \mathcal{A} to a sub- C^* -algebra of $C(\mathbb{T}, \mathcal{A})$, the C^* -algebra of continuous \mathcal{A} -valued functions on the unit circle; $S(\mathcal{A})$ consists of those functions $f : \mathbb{T} \rightarrow \mathcal{A}$ such that $f(1) = 0$ (or alternatively, $S(\mathcal{A}) = \mathcal{A} \otimes C_0(\mathbb{R})$, the tensor product of \mathcal{A} with the continuous, complex-valued functions on \mathbb{R} vanishing at infinity). A $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ yields the $*$ -homomorphism $S(\phi)$ defined by $([S\phi](f))(t) \mapsto \phi(f(t))$ for all $t \in \mathbb{T}$.

Note that this suspension functor is an extension of the one defined on topological spaces in the sense that for algebras $\mathcal{A} = C(X)$, we have that $S(\mathcal{A}) = C(S(X))$.

Definition 6.2.5. Operator K -theory is the sequence of functors K_n defined by $K_n = K_0(S^{|n|}(\mathcal{A}))$.

Generalized Bott periodicity provides natural isomorphisms $K_n \simeq K_{n+2}$. We are left with K_0 and $K_1 = K_0 \circ S$.

6.2.1 Unitalization

So far, we have defined operator K -theory only for the unital case. We describe the extension of K_0 to all C^* -algebras; a process we will be replicating later.

Definition 6.2.6. The minimal unitalization of a C^* -algebra \mathcal{A} (which itself may be unital or non-unital), is defined as the unital C^* -algebra \mathcal{A}^+ with underlying set $\mathcal{A} \times \mathbb{C}$, componentwise addition and scalar multiplication, and multiplication and involution given by

$$(a, z)(a', z') = (aa' + z'a + za', zz'), \quad (a, z)^* = (a^*, \bar{z}).$$

There exists a unique C^* -norm on \mathcal{A}^+ which we omit.

Note that $(-)^+$ is a functor from the category of C^* -algebras with $*$ -homomorphisms

to the category of unital C^* -algebras and unital $*$ -homomorphisms: a $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ yields $(a, z) \mapsto (\phi(a), z)$.

A copy of \mathcal{A} lives canonically inside \mathcal{A}^+ in the first component. Indeed, the unitalization of a C^* -algebra yields a short exact sequence

$$0 \rightarrow \mathcal{A} \xrightarrow{\iota} \mathcal{A}^+ \xrightarrow{\pi} \mathbb{C} \rightarrow 0$$

with ι being the injection into the first component and π being the projection to the second component. Exactness justifies identifying \mathcal{A} with $\ker \pi$.

Definition 6.2.7. For a C^* -algebra \mathcal{A} (either unital or non-unital) the K_0 group of \mathcal{A} is defined as the subgroup of $K_0(\mathcal{A}^+)$ given by the kernel of $K_0(\pi^+)$. A $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ yields a homomorphism between the kernels of $K_0(\mathcal{A}^+ \xrightarrow{\pi} \mathbb{C})$ and $K_0(\mathcal{B}^+ \xrightarrow{\pi} \mathbb{C})$ by restriction of $K_0(\phi^+)$ to the kernel of $K_0(\mathcal{A}^+ \xrightarrow{\pi} \mathbb{C})$.

6.2.2 Stability

Definition 6.2.8. The compact operators \mathcal{K} is the sub- C^* -algebra of $\mathcal{B}(\mathcal{H})$, with \mathcal{H} a Hilbert space of countable dimension, which is generated by the finite rank operators.

Alternatively, it is defined as the direct limit in the category of C^* -algebras of the sequence of matrix algebras

$$M_1(\mathbb{C}) \hookrightarrow M_2(\mathbb{C}) \hookrightarrow M_3(\mathbb{C}) \hookrightarrow \dots$$

where the injections are inclusion into the upper left corner: $x \mapsto \begin{pmatrix} x & 0 \\ 0 & 0 \end{pmatrix}$.

The C^* -algebra \mathcal{K} is *nuclear*, which means that, for any C^* -algebra \mathcal{A} there is a unique C^* -norm on the algebraic tensor product $\mathcal{A} \otimes_{alg} \mathcal{K}$ and thus we may speak unambiguously of the C^* -algebra $\mathcal{A} \otimes \mathcal{K}$.

Definition 6.2.9. The stabilization functor, which we (by a harmless abuse of notation) denote by \mathcal{K} , maps the category of C^* -algebras to itself by sending a C^* -algebra \mathcal{A} to the C^* -algebra $\mathcal{K}(\mathcal{A}) = \mathcal{A} \otimes \mathcal{K}$. It sends a $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ to $\mathcal{K}(\phi) : \mathcal{K}(\mathcal{A}) \rightarrow \mathcal{K}(\mathcal{B})$ defined by $\phi \otimes id_{\mathcal{K}}$.

It is alternatively defined as the direct limit of matrix algebras. That is, $\mathcal{K}(\mathcal{A})$ is the limit in the category of C^* -algebras of

$$M_1(\mathcal{A}) \hookrightarrow M_2(\mathcal{A}) \hookrightarrow M_3(\mathcal{A}) \hookrightarrow \dots$$

where the morphisms are inclusion into the upper left corner. A $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ yields homomorphisms $M_n(\phi) : M_n(\mathcal{A}) \rightarrow M_n(\mathcal{B})$ which form the components of the natural transformation yielding $\mathcal{K}(\phi)$.

The stabilization functor is an idempotent operation in the sense that $\mathcal{K} \circ \mathcal{K} \simeq \mathcal{K}$. This follows from the fact that $\mathcal{K} \otimes \mathcal{K} \simeq \mathcal{K}$.

Definition 6.2.10. A C^* -algebra \mathcal{A} is called *stable* or a *stabilization* if it is fixed (up to isomorphism) by the \mathcal{K} functor, i.e. $\mathcal{A} \simeq \mathcal{A} \otimes \mathcal{K}$.

Note that no stable C^* -algebra can be unital. Two C^* -algebras \mathcal{A} and \mathcal{B} are *stably equivalent* when $\mathcal{K}(\mathcal{A}) \simeq \mathcal{K}(\mathcal{B})$. Among stable C^* -algebras, stable equivalence reduces to ordinary isomorphism equivalence. As we shall see, stable algebras, in a certain sense, form the class of C^* -algebras for which the operator K -theory can be taken.

Theorem 6.2.11. Operator K -theory is matrix stable. That is, $K_0 \simeq K_0 \circ M_n$ and $K_1 \simeq K_1 \circ M_n$.

Theorem 6.2.12. Operator K -theory is continuous. That is, if

$$\mathcal{A}_1 \rightarrow \mathcal{A}_2 \rightarrow \mathcal{A}_3 \rightarrow \dots$$

is a direct sequence of C^* -algebras and $*$ -homomorphisms,

$$K_0(\mathcal{A}_1) \rightarrow K_0(\mathcal{A}_2) \rightarrow K_0(\mathcal{A}_3) \rightarrow \dots$$

is its image under the K_0 functor, and $\mathcal{A} = \varinjlim \mathcal{A}_n$, then,

$$K_0(\mathcal{A}) \simeq \varinjlim K_0(\mathcal{A}_n)$$

via the obvious homomorphism induced between cones. A similar statement holds for K_1 .

As a consequence of the preceding two theorems, and the alternative definition of the compact operators as the limit of a direct sequence of matrix algebras, we obtain:

Theorem 6.2.13. *Operator K -theory is stable. That is, $K_0 \simeq K_0 \circ \mathcal{K}$ and $K_1 \simeq K_1 \circ \mathcal{K}$.*

The operator K -theory functors are determined by their restrictions to stable C^* -algebras.

6.3 Alternative definition of operator K_0 functor

For a unital C^* -algebra \mathcal{A} , the K_0 -group can be expressed in a rather simple fashion in terms of projections of its stabilization [60, p105]. We will require this definition in the next section and thus describe it in explicit detail.

Definition 6.3.1. *Two projections P and Q in a C^* -algebra \mathcal{A} are unitarily equivalent if and only if there exists a unitary $u \in \mathcal{A}^+$ such that $P = uQu^*$.*

We denote unitary equivalence of P and Q in this subsection by $P \sim Q$.

Definition 6.3.2 (The Murray-von Neumann semigroup for unital \mathcal{A} , alternative definition). *The elements of the abelian monoid $V(\mathcal{A})$ are unitary equivalence classes of projections in $\mathcal{K}(\mathcal{A})$.*

The abelian addition is given by orthogonal addition. That is, if P and P' are two projections in $\mathcal{A} \otimes \mathcal{K}$, there exist projections Q and Q' such that $P \sim Q$, $P' \sim Q'$, and Q is orthogonal to Q' . The abelian addition operation on projections is defined by $[P] + [P'] = [Q + Q']$.

A $$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ yields an abelian monoid homomorphism by $[P] \mapsto [\mathcal{K}(\phi)(P)]$.*

By redefining the Murray-von Neumann semigroup functor V , we automatically get a new description of K_0 by composition with the Grothendieck group functor: $K_0 \simeq G \circ V$. Then, $K_0(\mathcal{A})$ is simply the collection of formal differences

$$[P] - [Q]$$

of elements of $V(\mathcal{A})$ with

$$[P] - [Q] = [P'] - [Q']$$

precisely when there exists $[R]$ such that

$$[P] + [Q'] + [R] = [P'] + [Q] + [R] .$$

Composing the action on morphisms of the Grothendieck group functor after the action of V just defined, we find that a unital $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ between unital C^* -algebras yields an abelian group homomorphism between the K_0 groups of \mathcal{A} and \mathcal{B} given by

$$[P] - [Q] \mapsto [\phi(P)] - [\phi(Q)] .$$

6.4 The extension of topological K -theory

The spatial diagram $G(\mathcal{A})$ associated to a unital C^* -algebra \mathcal{A} can be thought of as a geometric object which is a sort of topological approximation to the metaphorical noncommutative space underlying \mathcal{A} . Functors which are defined to act only on topological spaces can also be applied directly to $G(\mathcal{A})$ and the results compared with the generalization to noncommutative topology of the concepts they encapsulate (when they exist).

One might naively conjecture that by decomposing a noncommutative space into all its quotient spaces, taking the topological K -theory of all those quotient spaces which turn out to be genuine topological spaces, and pasting the results all together, one could recover the operator K -theory of the noncommutative space. Indeed, we can weaken this guess to hold only for stable C^* -algebras and not lose the ability to define K_0 .

Conjecture 6.4.1. $\tilde{K} \circ \mathcal{K}$ is naturally equivalent to $K_0 \simeq K_0 \circ \mathcal{K}$

Expanding the definition of \tilde{K} , as given in Definition 5.2.1, we see that the conjecture is for the following equation to hold:

Conjecture 6.4.2. $K_0 \simeq K_0 \circ \mathcal{K} \simeq \varinjlim K \circ \Sigma \circ g \circ \mathcal{K}$

In fact, we find that we must modify this conjecture in order to reconstruct operator K -theory as a colimit of the topological K -theory of its commutative subalgebras:

Theorem 6.4.3. $K_0 \simeq \varinjlim K \circ \Sigma \circ g_f \circ \mathcal{K}$

We have modified our conjecture by replacing g with g_f as defined in Definition 4.3.5.

Although it remains open whether $K_0 \simeq \tilde{K}$ we find that our original aim of reconstructing K_0 by applying K to spatial diagrams can be achieved by applying K not to the diagrams G but to subdiagrams $G_f = \Sigma \circ g_f$ after stabilizing.

Theorem 6.4.4. $K_0 \simeq K_0 \circ \mathcal{K} \simeq \tilde{K}_f \circ \mathcal{K}$ as functors from unital C^* -algebras to abelian groups.

Here, \tilde{K}_f is defined for unital C^* -algebras in a manner most similar to how \tilde{K} is: as $\varinjlim K \circ G_f$. As all C^* -algebras isomorphic to $\mathcal{A} \otimes \mathcal{K}$ lack a unit, we must extend \tilde{K}_f to all C^* -algebras and we do so using the same method used for K_0 : if $\pi : \mathcal{A}^+ \rightarrow \mathbb{C}$ is the projection from the unitalization of \mathcal{A} to the sub- C^* -algebra of scalar multiples of the unit, then $\tilde{K}_f(\mathcal{A})$ is defined as the kernel of $\tilde{K}_f(\pi)$ [26, p186].

Proof. For unital \mathcal{A} , we will compute

$$\tilde{K}_f \circ \mathcal{K}(\mathcal{A}) = \varinjlim K \circ G_f(\mathcal{A} \otimes \mathcal{K})$$

in stages and present it in such a form which makes it clear that it is isomorphic in a natural way to $K_0(\mathcal{A})$. We begin by describing how the functor \tilde{K}_f acts on unital algebras before generalizing to the non-unital case.

The objects of $S_f(\mathcal{A})$ are the unital, finite-dimensional, commutative sub- C^* -algebras of \mathcal{A} . The morphisms are given by the restrictions of inner automorphisms. These morphisms are all of the form $i \circ r$ where i is an inclusion and r is an isomorphism between subalgebras which are related by unitary rotation.

Under the Gel'fand spectrum functor, the image of such an object is a finite, discrete space whose points are in correspondence with the atomic projections of the subalgebra. The images of the inclusions $i : U \rightarrow V$ are functions $\Sigma(i) : \Sigma(V) \rightarrow \Sigma(U)$ with the property that whenever a point $p \in \Sigma(U)$ corresponds to a projection P atomic in U , then P is the sum of the atomic projections in V which correspond

to the points of $[\Sigma(i)]^{-1}(p)$. The isomorphisms are sent by the spectrum functor to bijections which connect points corresponding to unitarily equivalent projections.

Under the topological K -functor, each object yields a direct sum of copies of \mathbb{Z} : one for each point (a trivial vector bundle of each dimension and formal inverses). Taking the colimit of the diagram then yields, as described in the previous chapter, a direct sum of the groups indexed by the objects of $S_f(\mathcal{A})$ modulo the relations generated by the morphisms of the diagram. In our case, this is a quotient of the direct sum of a copy of \mathbb{Z} for each pair (U, P) where U is a finite-dimensional, unital, commutative sub- C^* -algebra of \mathcal{A} and P is an atomic projection in U .

The image under $K \circ \Sigma$ of the inclusions result in the identifications of a generator of a copy of \mathbb{Z} associated to a pair (U, P) with the sums of generators associated to pairs (V, P_i) whenever $U \subset V$ and P_i sums to P . Every projection $P \in \mathcal{A}$ is an atomic projection in the subalgebra $\mathbb{C}P + \mathbb{C}(1 - P)$ which is included in every subalgebra which contains P as a member. As the generators associated to the same projection P atomic in different subalgebras are all identified in the colimit, we see that we may speak of the element of the colimit group $[(P)]$ associated to P without reference to which subalgebra it appears in. Thus, the abelian group $\tilde{K}_f(\mathcal{A})$ can be viewed as a quotient of the free abelian group generated by the elements $[(P)]$. The first class of identifications consists of those between elements $[(P)]$ with the sum of elements $[(P_i)]$ whenever P_i are mutually orthogonal and sum to P . The isomorphisms in the diagram provide the second class; they ensure that the elements associated to unitarily equivalent projections are also identified.

For non-unital algebras such as $\mathcal{A} \otimes \mathcal{K}$, we need to determine the kernel of $\tilde{K}_f(\pi)$ with π the canonical projection from $(\mathcal{A} \otimes \mathcal{K})^+$ to \mathbb{C} . This is not so difficult, however, as all projections in $(\mathcal{A} \otimes \mathcal{K})^+$ are of the form P or $1 - P$ for $P \in \mathcal{A} \otimes \mathcal{K}$. As

$$[(1 - P)] = [(1)] - [(P)]$$

we see that all elements of the colimit group can be expressed using only elements

associated either to the identity projection or to projections in $\mathcal{A} \otimes \mathcal{K}$. It is precisely those elements of $\tilde{K}_f((\mathcal{A} \otimes \mathcal{K})^+)$ which can be expressed using only elements associated to projections in $\mathcal{A} \otimes \mathcal{K}$ which are in the kernel.

We are ready to define the natural transformation $\eta : K_0 \rightarrow \tilde{K}_f \circ \mathcal{K}$. The component of the natural isomorphism $\eta_{\mathcal{A}}$ sends $[P] - [Q] \in K_0(\mathcal{A})$ to $[(P)] - [(Q)]$. This is easily seen to be well defined, for if

$$[P] + [Q'] + [R] = [P'] + [Q] + [R]$$

then we may find mutually orthogonal representatives of all these projections and demonstrate that

$$[(P)] - [(Q)] = [(P')] - [(Q')].$$

Preservation of addition follows by a similar argument.

We define an inverse map to demonstrate bijectivity. The element $[(P)] \in \tilde{K}_f(\mathcal{A} \otimes \mathcal{K})$ is sent by $\eta_{\mathcal{A}}^{-1}$ to $[P]$. Since the equivalences induced by the morphisms of the diagram are respected at the level of K -theory, this map is well-defined.

To demonstrate the naturality of these isomorphisms, let $\phi : \mathcal{A} \rightarrow \mathcal{B}$ be a unital $*$ -morphism. We will require the naturally induced $*$ -morphism $\mathcal{K}\phi : \mathcal{A} \otimes \mathcal{K} \rightarrow \mathcal{B} \otimes \mathcal{K}$ which is defined explicitly by $\phi \otimes id_{\mathcal{K}}$.

$$\begin{array}{ccc} K_0(\mathcal{A}) & \xrightarrow{\eta_{\mathcal{A}}} & \varinjlim \tilde{K}_f \circ \mathcal{K}(\mathcal{A}) \\ \downarrow K_0(\phi) & & \downarrow \varinjlim \tilde{K}_f \circ \mathcal{K}(\phi) \\ K_0(\mathcal{B}) & \xrightarrow{\eta_{\mathcal{B}}} & \varinjlim \tilde{K}_f \circ \mathcal{K}(\mathcal{B}) \end{array}$$

An element $[P] - [Q]$ in $K_0(\mathcal{A})$ is mapped to $[(P)] - [(Q)]$ by $\eta_{\mathcal{A}}$. This is then mapped to $[(\mathcal{K}\phi(P))] - [(\mathcal{K}\phi(Q))]$ by $\varinjlim \tilde{K}_f \circ \mathcal{K}(\phi)$. Alternatively, $[P] - [Q]$ in $K_0(\mathcal{A})$ is mapped to $[\mathcal{K}\phi(P)] - [\mathcal{K}\phi(Q)]$ in $K_0(\mathcal{B})$ which is in turn mapped to $[(\mathcal{K}\phi(P))] - [(\mathcal{K}\phi(Q))]$ by $\eta_{\mathcal{B}}$.

□

Chapter 7

Noncommutative Topology

It is conjectured that applying the extension process to the functor which assigns a compact Hausdorff space its lattice of open sets yields the functor which assigns to a unital C^* -algebra its complete meet-semilattice of closed, two-sided ideals. This is tantamount to taking the limit of topologies of each context to recover the hull-kernel topology on the C^* -algebra's primary ideal spectrum (the analogue of the spectrum of commutative ring theory). After reviewing the primary ideal spectrum and some necessary facts about ideals and projections, this conjecture is then rephrased as a characterization of those partial ideals which arise from total ideals. In joint work with Rui Soares Barbosa, the von Neumann algebraic version of the conjecture is proved.

7.1 The primitive ideal space

Here, we include some basic facts on prime ideals of rings and, its C^* -algebraic analogue, the primitive ideal space which are required in our explication of the motivation for the consideration of the extension of the closed-set lattice functor which follows.

7.1.1 The spectrum of commutative rings

In commutative ring theory and algebraic geometry, the starting point for the application of geometrical methods is the association of a topological space to rings [35, p70] (it is, in fact, a locally ringed space; however, we will not be considering this additional structure).

Definition 7.1.1. A prime ideal J of a commutative ring R is a ideal $J \subsetneq R$ such that whenever we have $a, b \in R$ such that $ab \in J$ then either $a \in J$ or $b \in J$.

The canonical examples of prime ideals come from the ideals of the ring of integers generated by prime numbers.

Definition 7.1.2. Let R be a commutative ring and let $I \subset R$ be a two-sided ideal of R . Then $\text{hull}(I)$ is the set of prime ideals containing I .

Definition 7.1.3. Let R be a commutative ring. The contravariant spectrum functor Spec from the category of rings and ring morphisms to the category of topological spaces and continuous maps sends an object R to the topological space whose points are the prime ideals of R . It is equipped with the hull-kernel (or Zariski, or Jacobson) topology whose closed sets are of the form $\text{hull}(I)$ for some two-sided ideal $I \subset R$.

The spectrum functor sends a ring morphism $h : R \rightarrow S$ to the continuous map $\text{Spec} : \text{Spec}(R) \rightarrow \text{Spec}(S)$ which sends a prime ideal J to its preimage $h^{-1}(J)$ under

h .

7.1.2 The primitive ideal space

These definitions and theorems can be found in [5, p208] and [12, p118].

Definition 7.1.4. A primitive ideal J of a C^* -algebra \mathcal{A} is one which is the kernel of an irreducible representation of \mathcal{A} .

Recall that an irreducible representation of a C^* -algebra \mathcal{A} is a $*$ -representation $\pi : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{H})$ such that for no $a \in \mathcal{A}$ and nontrivial closed subspaces $S \subset \mathcal{H}$ is it the case that $\pi(a)S \subset S$. Every pure state of \mathcal{A} gives rise to an irreducible representation \mathcal{A} by the Gel'fand-Naimark-Segal construction.

Definition 7.1.5. Let \mathcal{A} be a C^* -algebra and let $I \subset \mathcal{A}$ be a closed, two-sided ideal of \mathcal{A} . Then $\text{hull}(I)$ is the set of primitive ideals containing I .

Definition 7.1.6. Let \mathcal{A} be a C^* -algebra. The primitive ideal space $\text{Prim}(\mathcal{A})$ is the set of the primitive ideals of \mathcal{A} . It is equipped with the hull-kernel (or Zariski, or Jacobson) topology whose closed sets are of the form $\text{hull}(I)$ for some two-sided norm closed ideal $I \subset \mathcal{A}$.

Theorem 7.1.7. The map hull is an order preserving bijection between the set of two-sided ideals of a C^* -algebra \mathcal{A} and the closed sets of the hull-kernel topology on $\text{Prim}(\mathcal{A})$.

Definition 7.1.8. The spectrum $\hat{\mathcal{A}}$ of a C^* -algebra \mathcal{A} is the set of unitary equivalence classes of irreducible representations of \mathcal{A} . It is equipped with the coarsest topology with respect to which the map $[\pi] \rightarrow \ker \pi$ is continuous.

The topology on $\hat{\mathcal{A}}$ is thus also order isomorphic to the partially ordered set of two-sided ideals of \mathcal{A} .

7.2 Von Neumann Algebras

In this subsection, we briefly outline some elementary facts about von Neumann algebras [5, Chap. 3] which we will require in the rest of this section.

Definition 7.2.1. A unital von Neumann algebra \mathcal{A} is a unital $*$ -subalgebra of $\mathcal{B}(\mathcal{H})$ which is closed in the weak (operator) topology.

Recall that a net of operators (T_α) in $\mathcal{B}(\mathcal{H})$ converges to T in the *weak topology* if and only if, for every vector $v \in \mathcal{H}$ and linear functional $\phi \in \mathcal{H}^*$, we have that $(\phi(T_\alpha(v)))$ converges to $\phi(T(v))$. As convergence of a net of operators in norm implies its weak convergence, we see that von Neumann algebras are examples of C^* -algebras. We may equally well have defined von Neumann algebras to be $*$ -subalgebras of $\mathcal{B}(\mathcal{H})$ which are closed in the strong, ultraweak, or ultrastrong topologies as the closures of $*$ -subalgebras of $\mathcal{B}(\mathcal{H})$ in these topologies all coincide. Von Neumann proved that taking any of these closures of unital $*$ -subalgebras of $\mathcal{B}(\mathcal{H})$ coincides also with taking the double commutant (though he did not know of the ultrastrong topology).

We will primarily require facts about projections and ideals of von Neumann algebras and the relationship between the two notions.

7.2.1 Projections

The *projections* of \mathcal{A} are operators p such that $p = p^* = p^2$ and they are orthogonal projections onto closed subspaces of \mathcal{H} . This yields a natural ordering on projections induced by the inclusion relation on their corresponding subspaces. Alternatively, the ordering can be defined by:

$$p \leq q \text{ if and only if } pq = p \text{ if and only if } qp = p .$$

We denote the partially ordered set of projections in \mathcal{A} by $\mathcal{P}(\mathcal{A})$. In von Neumann algebras, the collection of projections forms a complete lattice: the infimum $\inf_{\alpha} p_{\alpha}$ of an arbitrary collection of projections $\{p_{\alpha}\}$ is given by the orthogonal projection onto $\bigcap p_{\alpha} \mathcal{H}$ whereas the supremum $\sup_{\alpha} p_{\alpha}$ is the orthogonal projection onto the closed linear span of $\bigcup p_{\alpha} \mathcal{H}$. The orthogonal complement map \perp which sends p to $1 - p$ makes this lattice complemented in the sense that $p \vee p^{\perp} = 1$, $p \wedge p^{\perp} = 0$, and $p^{\perp\perp} = p$.

The set of projections in \mathcal{A} is also equipped with several other partial orderings which arise from the canonical partial ordering and certain compatible equivalence relations. We will require, in particular, the notions of Murray-von Neumann equivalence of projections and unitary equivalence of projections.

The intuition behind Murray-von Neumann equivalence is to identify projections whose corresponding image subspaces are of the same dimension. That is, there should be an operator $v \in \mathcal{A}$ mapping the Hilbert space \mathcal{H} to itself which isometrically maps the subspace of one projection to the subspace of another, thereby witnessing the equality of their dimension.

Definition 7.2.2. *Two projections p and q in a von Neumann algebra \mathcal{A} are Murray-von Neumann equivalent, denoted $p \sim q$, if and only if there exists $v \in \mathcal{A}$ such that*

$$p = v^*v \text{ and } q = vv^*.$$

The partial ordering on $\mathcal{P}(\mathcal{A})$ induces a partial ordering on the set of Murray-von Neumann equivalence classes of projections. We write $p \preceq_{\mathbf{M}} q$ to denote that $p \sim p'$ for some $p' \leq q$.

Definition 7.2.3. *Two projections p and q in a unital von Neumann algebra \mathcal{A} are unitarily equivalent, denoted $p \sim_u q$, if and only if there exists a unitary element $u \in \mathcal{A}$ such that $p = uqu^*$.*

Unitary equivalence is stronger than Murray-von Neumann equivalence. The partial ordering on $\mathcal{P}(\mathcal{A})$ induces a partial ordering on the set of unitary equivalence classes of projections. We write $p \preceq_u q$ to denote that $p \sim_u p'$ for some $p' \leq q$.

We will require the fact that when p and q are orthogonal and Murray-von Neumann equivalent they are automatically also unitarily invariant [6, p445].

Definition 7.2.4. *The central carrier $C(p)$ of a projection $p \in \mathcal{A}$ is the projection*

$$\inf\{z \in \mathcal{P}(\mathcal{A}) \cap \mathcal{Z}(\mathcal{A}), z \geq p\}.$$

It is immediate from this definition that a projection p and a unitary rotation upu^* have the same central carrier for $z \geq p$ if and only if $uzu^* = zuu^* = z \geq upu^*$. It is also immediate that if $S \subset \mathcal{P}(\mathcal{A})$ is a set of projections, then $C(\sup S) = \sup_{p \in S} C(p)$.

One of the basic technical tools we will require is the comparison theorem of projections in a von Neumann algebra [44]. The intuitive idea is best understood in a factor (a von Neumann algebra with trivial centre) which can be thought of as an elementary direct summand. Here, the dimension of two projections can be compared; either they are of equal dimension, or the dimension of one exceeds the dimension of the other.

Theorem 7.2.5 (Comparison theorem). *Let p and q be projections in $\mathcal{P}(\mathcal{A})$. There exists a central projection $z \in \mathcal{A}$ such that*

$$zp \succ_{\mathbf{M}} zq \quad \text{and} \quad z^\perp p \prec_{\mathbf{M}} z^\perp q.$$

7.2.2 Ideals

Ideals of operator algebras must satisfy both the usual algebraic conditions as well as an additional topological condition.

It turns out that the appropriate notion of morphism for von Neumann algebras is not weakly continuous $*$ -homomorphism but rather ultraweakly continuous $*$ -homomorphism. The ultraweak topology is stronger than the weak topology.

Definition 7.2.6. A left (resp. right) ideal I of a von Neumann algebra \mathcal{A} is a left (resp. right) ring ideal $I \subset \mathcal{A}$ that is closed in the ultraweak topology.

A total ideal or two-sided ideal I of a von Neumann algebra \mathcal{A} is two-sided ring ideal $I \subset \mathcal{A}$ that is closed in the ultraweak topology.

Left, right, and total ideals correspond with projections. Examples of left (resp. right) ideals are the sets given by $\mathcal{A}p$ (resp. $p\mathcal{A}$); these are the kernels of morphisms given by right (resp. left) multiplication by p^\perp .

Theorem 7.2.7. Every left ideal $L \subset \mathcal{A}$ of a von Neumann algebra \mathcal{A} is of the form $L = \mathcal{A}p$ for a projection $p \in \mathcal{P}(\mathcal{A})$. Further, the projection p is uniquely determined by L .

Total ideals are precisely those left or right ideals corresponding to central projections.

Theorem 7.2.8. Every total ideal $I \subset \mathcal{A}$ of a von Neumann algebra \mathcal{A} is of the form $I = z\mathcal{A}z = z\mathcal{A} = \mathcal{A}z$ for a unique central projection $z \in \mathcal{P}(\mathcal{A}) \cap \mathcal{Z}(\mathcal{A})$.

7.3 Total and partial ideals

In the previous chapter, we found that including the class of restricted automorphisms in the diagram of topological spaces meant to act as a noncommutative spectrum allowed us to give a novel definition of operator K -theory using nearly the precise method which we conjectured in Chapter 2.

The next step in using extensions to directly obtain noncommutative analogues from basic topological concepts would be to establish the conjecture that extending the notion of closed topological subset leads to the analogous algebraic concept of closed, two-sided ideal.

The principal theorem proved in this chapter (Theorem 7.3.6) concerns the von Neumann algebraic analogue of a conjecture made by the author for C^* -algebras; its proof is joint work with Rui Soares Barbosa. A notion of *partial ideal* for an operator algebra is a weakening the notion of ideal where the defining algebraic conditions are enforced only in the commutative subalgebras. We show that, in a von Neumann algebra, the ultraweakly closed two-sided ideals, which we call *total ideals*, correspond to the unitarily invariant partial ideals. The result also admits an equivalent formulation in terms of central projections. We describe the original question concerning C^* -algebras before stating and proving the von Neumann algebraic version.

7.3.1 Motivation


In this subsection, we motivate the characterization of total ideals as invariant partial ideals.

It is conjectured that taking the extension of the notion of open subset of a space (or equivalently, closed subset) would lead to the notion of closed, two-sided (i.e.

total) ideal of a C^* -algebra. To formalize this idea, let $\mathcal{T} : \mathbf{KHaus} \rightarrow \mathbf{CMSLat}$ be the functor which assigns to a compact Hausdorff topological space its complete lattice of closed sets (with $C_1 \leq C_2$ if and only if $C_1 \supset C_2$) and assigns to a continuous function the complete meet-semilattice homomorphism mapping a closed set to its image under the continuous function and let $\tilde{\mathcal{T}}$ be its extension. Suppose further that \mathcal{I} is the contravariant functor from the category of C^* -algebras to the category of complete meet-semilattices which sends an algebra to its lattice of total ideals and a $*$ -homomorphism $\phi : \mathcal{A} \rightarrow \mathcal{B}$ to the homomorphism of complete meet-semilattices $\mathcal{I}(\phi) : \mathcal{I}(\mathcal{B}) \rightarrow \mathcal{I}(\mathcal{A})$ taking an ideal $I \subset \mathcal{B}$ to the ideal $\phi^{-1}(I)$ of \mathcal{A} . Note that \mathcal{I} , once restricted to the full subcategory of commutative C^* -algebras, is naturally equivalent to $\mathcal{T} \circ \Sigma$.

Conjecture 7.3.1. *The functors $\tilde{\mathcal{T}}$ and \mathcal{I} are naturally isomorphic.*

To prove Conjecture 7.3.1 would be essentially to demonstrate a bijective correspondence of total ideals of \mathcal{A} with certain functions π which map commutative sub- C^* -algebras V of \mathcal{A} to closed ideals of V . To see this, note that the limit lattice $\tilde{\mathcal{T}}(\mathcal{A}) = \varprojlim \mathcal{T}G(\mathcal{A})$ is a cone over the diagram $\mathcal{T}G(\mathcal{A})$.



As we shall soon see, the elements of $\tilde{\mathcal{T}}(\mathcal{A})$ are precisely what we call the *invariant partial ideals* of \mathcal{A} , i.e. choices of elements from each $\mathcal{I}(V)$ subject to the condition of the remark following Definition 7.3.4.

Proving Conjecture 7.3.1 would establish a strong relationship between the topologies of the geometric object $G(\mathcal{A})$ and $\text{Prim}(\mathcal{A})$, the primitive ideal space of \mathcal{A} : we would be able to recover the lattice of the hull-kernel topology on $\text{Prim}(\mathcal{A})$, as the limit of the topological lattices of the object $G(\mathcal{A})$. Establishing this conjecture would allow considering G to be an enrichment of Prim . Given that Prim is a C^* -algebraic variant of the ring-theoretic spectrum whose hull-kernel topology provides the basis for sheaf-theoretic ring theory, one might hope to speculate about the use of sheaf-theoretic methods in noncommutative topology.

7.3.2 Partial and total ideals of C^* -algebras

All algebras and subalgebras considered throughout this chapter are assumed to be unital. By a *total ideal* of a C^* -algebra \mathcal{A} , we mean a norm closed, two-sided ideal of \mathcal{A} .

Definition 7.3.2. A partial ideal of a C^* -algebra \mathcal{A} is a map π that assigns to each commutative sub- C^* -algebra V of \mathcal{A} a closed ideal of V such that $\pi(V) = \pi(V') \cap V$ whenever $V \subset V'$.

Remark 1. A partial ideal is precisely a choice of $\pi(V) \in \mathcal{I}(V)$ for each commutative sub- C^* -algebra V of \mathcal{A} such that whenever there is an inclusion morphism $\iota : V \hookrightarrow V'$, then

$$\pi(V) = \mathcal{I}(\iota)(\pi(V')) = \pi(V') \cap V ;$$

i.e. the following diagram commutes.

$$\begin{array}{ccc}
 & & \{*\} \xrightarrow{* \mapsto \pi(V')} \mathcal{I}(V') \\
 & \uparrow \iota & \searrow \downarrow \mathcal{I}(\iota) \\
 V' & & \{*\} \xrightarrow{* \mapsto \pi(V)} \mathcal{I}(V)
 \end{array} \tag{7.1}$$

Remark 2. *The concept of partial ideal was introduced by Reyes [58] in the more general context of partial C^* -algebras. His definition differs slightly but is equivalent in our case: a subset P of normal elements of \mathcal{A} such that $P \cap V$ is a closed ideal of V for all commutative sub- C^* -algebras V of \mathcal{A} .*

Partial ideals exist in abundance: every closed, left (or right) ideal I of \mathcal{A} gives rise to a partial ideal π_I in a natural way by choosing $\pi_I(V)$ to be $I \cap V$.

For example, in a matrix algebra $M_n(\mathbb{C})$, the right ideal $pM_n(\mathbb{C})$, for $p \in M_n(\mathbb{C})$ a nontrivial projection, yields a nontrivial partial ideal of $M_n(\mathbb{C})$ in this way. As matrix algebras are simple, it cannot be the case that these nontrivial partial ideals also arise as π_I from a total ideal I . This raises a natural question:

Question 7.3.3. *Which partial ideals of C^* -algebras arise from total ideals?*

Some partial ideals do not even arise from left or right ideals: for example, choosing arbitrary nontrivial ideals from every nontrivial commutative sub- C^* -algebra of $M_2(\mathbb{C})$ yields, in nearly all cases, nontrivial partial ideals of $M_2(\mathbb{C})$. However, a hint towards identifying those partial ideals which arise from total ideals is given by a simple observation. If $\text{Ad}_u : \mathcal{A} \rightarrow \mathcal{A}$ is an inner automorphism of \mathcal{A} —that is, one given by conjugation by a unitary u of \mathcal{A} —then $\text{Ad}_u(I) = I$ for any total ideal $I \subset \mathcal{A}$. This imposes a special condition on the partial ideal $\pi_I(V) = I \cap V$ which arises from I .

Definition 7.3.4. *An invariant partial ideal π of a C^* -algebra \mathcal{A} is a partial ideal of \mathcal{A} such that, for each commutative sub- C^* -algebra $V \subset \mathcal{A}$ and any unitary $u \in \mathcal{A}$, the rotation by u of the ideal associated to V is the ideal associated to the rotation by u of V . That is,*

$$\text{Ad}_u(\pi(V)) = \pi(\text{Ad}_u(V)) .$$

Remark 3. *Imposing the invariance condition on partial ideals is equivalent to extending the requirement on maps π of Diagram (7.1) from inclusion maps to all*

$*$ -homomorphisms $\text{Ad}_u|_V^{V'} : V \longrightarrow V'$ arising as a restriction of the domain and codomain of an inner automorphism. An invariant partial ideal is precisely a choice of $\pi(V) \in \mathcal{I}(V)$ for each commutative sub- C^* -algebra V of \mathcal{A} such that whenever there is a morphism $\text{Ad}_u|_V^{V'} : V \longrightarrow V'$ as above, then

$$\pi(V) = \mathcal{I}(\text{Ad}_u|_V^{V'}) (\pi(V')) = \text{Ad}_{u*} (\pi(V')) \cap V ;$$

i.e. the following diagram commutes.

$$\begin{array}{ccc} V' & & \{*\} \xrightarrow{* \mapsto \pi(V')} \mathcal{I}(V') \\ \uparrow \text{Ad}_u|_V^{V'} & & \searrow * \mapsto \pi(V) \quad \downarrow \mathcal{I}(\text{Ad}_u|_V^{V'}) \\ V & & \mathcal{I}(V) \end{array}$$

Thus, we arrive at the simplest possible conjecture:

Conjecture 7.3.5. *A partial ideal of a C^* -algebra arises from a total ideal if and only if it is an invariant partial ideal. Consequently, the map $I \longmapsto \pi_I$ is a bijective correspondence between total ideals and invariant partial ideals.*

Note that the first part of the statement says that the map $I \longmapsto \pi_I$ is surjective onto the invariant partial ideals. The second part of the statement follows easily from this, since injectivity of this map is obvious: the left inverse is given by mapping an invariant partial ideal of the form π_I to the linear span of $\bigcup_V \pi(V)$, which is equal to I itself.

7.3.3 Partial and total ideals of von Neumann algebras

One may define *partial ideal* (resp. *invariant partial ideal*) for a von Neumann algebra by replacing in Definition 7.3.2 (resp. Definition 7.3.4) the occurrences of “commutative sub- C^* -algebra” with “commutative sub-von-Neumann-algebra” and

“closed ideal” with “ultraweakly closed ideal”. A total ideal of a von Neumann algebra is, as in Definition 7.2.6, an ultraweakly closed, two-sided ideal. As before, a total ideal I determines an invariant partial ideal π_I in the same way, and the map $I \mapsto \pi_I$ is injective.

Establishing the analogue of Conjecture 7.3.5 for von Neumann algebras provides some measure of evidence for the original conjecture’s verity, and its proof may be adapted to show that the original conjecture holds for a large class of—or perhaps all— C^* -algebras.

Theorem 7.3.6 (Principal theorem of chapter). *A partial ideal of a von Neumann algebra arises from a total ideal if and only if it is an invariant partial ideal. Consequently, the map $I \mapsto \pi_I$ is a bijective correspondence between total ideals and invariant partial ideals.*

Total ideals of a von Neumann algebra \mathcal{A} are in bijective correspondence with central projections z of \mathcal{A} : every total ideal I is of the form $z\mathcal{A}$ for a unique z (Theorem 7.2.8) [5]. This allows us to rephrase the theorem in terms of projections which are vastly more convenient to work with.

Definition 7.3.7. *A consistent family of projections of a von Neumann algebra \mathcal{A} is a map Φ that assigns to each commutative sub-von-Neumann-algebra V of \mathcal{A} a projection in V such that:*

1. *for any V and V' such that $V \subset V'$, $\Phi(V)$ is the largest projection in V which is less than or equal to $\Phi(V')$, i.e.*

$$\Phi(V) = \sup \{q \text{ is a projection in } V \mid q \leq \Phi(V')\} .$$

An invariant family of projections is such a map which further satisfies

2. *for any unitary element $u \in \mathcal{A}$, $\Phi(uVu^*) = u\Phi(V)u^*$.*

The correspondence between total ideals and central projections yields correspondences between partial ideals (resp. invariant partial ideals) and consistent (resp. invariant) families of projections. We therefore establish Theorem 7.3.6 in the third section by proving an equivalent statement. Just as was the case for ideals, any projection p determines a consistent family of projections Φ_p defined by choosing $\Phi_p(V)$ to be the largest projection p in V which is less than or equal to p . For a central projection z , Φ_z turns out to be an invariant family. In the opposite direction, any consistent family of projections Φ gives a central projection $\Phi(\mathcal{Z}(\mathcal{A}))$ where $\mathcal{Z}(\mathcal{A})$ is the centre of \mathcal{A} .

Theorem 7.3.8 (Principal theorem of chapter, reformulated). *A consistent family of projections of a von Neumann algebra arises from a central projection if and only if it is an invariant family of projections. Consequently, the maps $z \mapsto \Phi_z$ and $\Phi \mapsto \Phi(\mathcal{Z}(\mathcal{A}))$ define a bijective correspondence between central projections and invariant families of projections.*

7.4 Technical Preliminaries

7.4.1 Little lemmata

In proving our main result, we shall make use of some simple properties of consistent families of projections which we record here as lemmata for clarity.

Lemma 7.4.1. *Let \mathcal{A} be a von Neumann algebra, and Φ be a consistent family of projections. Suppose $V \subset V'$ are commutative sub-von-Neumann-algebras of \mathcal{A} . Then:*

- i. $\Phi(V) \leq \Phi(V')$;
- ii. if $p \in V$ and $p \leq \Phi(V')$, then $p \leq \Phi(V)$;
- iii. in particular, if $\Phi(V') \in V$, then $\Phi(V') = \Phi(V)$.

Proof. Properties i and ii are simple consequences of the requirement in the definition of consistent family of projections that $\Phi(V)$ is the largest projection in V smaller than $\Phi(V')$. Property iii is a particular case of ii. \square

Given a commutative subset X of \mathcal{A} , denote by V_X the commutative sub-von-Neumann-algebra of \mathcal{A} generated by X and the centre $\mathcal{Z}(\mathcal{A})$, i.e. $V_X = (X \cup \mathcal{Z}(\mathcal{A}))''$. Given a finite commutative set of projections $\{p_1, \dots, p_n\}$, we write V_{p_1, \dots, p_n} for $V_{\{p_1, \dots, p_n\}}$.

Lemma 7.4.2. *Let \mathcal{A} be a von Neumann algebra; Φ a consistent family of projections; M a commutative set of projections in \mathcal{A} such that $\Phi(V_m) \geq m$ for all $m \in M$; and s the supremum of the projections in M . Then $\Phi(V_s) \geq s$.*

Proof. For all $m \in M$, since $V_m \subseteq V_M$, we have

$$\Phi(V_M) \geq \Phi(V_m) \geq m$$

by Lemma 7.4.1-i and the assumption that $\Phi(V_m) \geq m$. Hence, $\Phi(V_M)$ is at least the supremum of the projections in M , i.e. $\Phi(V_M) \geq s$. Now, from $V_s \subset V_M$ and $s \in V_s$, we conclude by Lemma 7.4.1-ii that $s \leq \Phi(V_s)$. \square

7.4.2 Partial orthogonality

Definition 7.4.3. Two projections p and q are partially orthogonal whenever there exists a central projection z such that zp and zq are orthogonal while $z^\perp p$ and $z^\perp q$ are equal.

Note that partially orthogonal projections necessarily commute and that if p_1 and p_2 are partially orthogonal, so is the pair zp_1 and zp_2 for any central projection z . A set of projections is *partially orthogonal* whenever any pair of projections in the set is partially orthogonal. We will require in the sequel the following simple lemma:

Lemma 7.4.4. Let p_1 and p_2 be projections and z be a central projection such that zp_1 and zp_2 are partially orthogonal and $z^\perp p_1$ and $z^\perp p_2$ are partially orthogonal. Then p_1 and p_2 are partially orthogonal.

Proof. As zp_1 and zp_2 are partially orthogonal, there exists a central projections y such that

$$yzp_1 = yzp_2 \quad \text{and} \quad y^\perp zp_1 \perp y^\perp zp_2 .$$

Similarly, as $z^\perp p_1$ and $z^\perp p_2$ are partially orthogonal, there exists a central projections x such that

$$xz^\perp p_1 = xz^\perp p_2 \quad \text{and} \quad x^\perp z^\perp p_1 \perp x^\perp z^\perp p_2 .$$

Summing both statements above, we conclude that

$$(yz + xz^\perp)p_1 = (yz + xz^\perp)p_2 \quad \text{and} \quad (y^\perp z + x^\perp z^\perp)p_1 \perp (y^\perp z + x^\perp z^\perp)p_2 ,$$

where $yz + xz^\perp$ is a central projection and $(yz + xz^\perp)^\perp = y^\perp z + x^\perp z^\perp$. So, p_1 and p_2 are partially orthogonal. \square

7.4.3 Main lemma

When comparing projections, we write \leq to denote the usual order on projections, \preceq_M for the order up to Murray-von Neumann equivalence, and \preceq_u for the order up to unitary equivalence.

The following lemma is one of the main steps of the proof. The idea is to start with a projection q in a von Neumann algebra and to cover, as much as possible, its central carrier $C(q)$ by a commutative subset of the unitary orbit of q . The lemma states that, in order to cover $C(q)$ with projections from the unitary orbit of q , it suffices to take a commutative subset, M , and (at most) one other projection, uqu^* , which is strictly larger than the remainder $C(q) - \sup M$. That is, the remainder from what can be covered by a commutative set is strictly smaller than q up to unitary equivalence.

Lemma 7.4.5. *Let q be a projection in a von Neumann algebra \mathcal{A} . Then there exists a set M of projections such that:*

- i. $q \in M$;
- ii. M is a subset of the unitary orbit of q ;
- iii. M is a commutative set;
- iv. the supremum s of M satisfies

$$s^R \prec_u q$$

where $s^R = C(q) - s$.

Proof. Let O be the unitary orbit of q . The partially orthogonal subsets of O which contain q form a poset under inclusion. Given a chain in this poset, its union is partially orthogonal: any two projections in the union must appear together somewhere in one subset in the chain and are thus partially orthogonal. Hence, by Zorn's lemma, we can construct a maximal partially orthogonal subset M of the unitary orbit of q such that $q \in M$. Clearly, M satisfies conditions i–iii.

Denote by s the supremum of the projections in M . Its central carrier $C(s)$ is equal to the central carrier $C(q)$ of q . This is because $C(-)$ is constant on unitary orbits and $C(\sup_{m \in M} m) = \sup_{m \in M} C(m)$. We now need to show that $s^R \prec_u q$.

By the comparison theorem for projections in a von Neumann algebra (Theorem 7.2.5), there is a central projection z such that

$$zs^R \succ_M zq \quad \text{and} \quad z^\perp s^R \prec_M z^\perp q.$$

We can assume without loss of generality that $z \leq C(q)$ since

$$C(q)^\perp q = C(q)^\perp s^R = 0.$$

Moreover, as s and s^R are orthogonal, there is a unitary which witnesses these order relationships. That is, there is a unitary u such that

$$zs^R \geq z(uku^*) \quad \text{and} \quad z^\perp s^R < z^\perp(uku^*).$$

We will show that z vanishes and thus conclude that $s^R < uku^*$ as required.

Define v to be the unitary $zu + z^\perp 1$ which acts as u within the range of z and as the identity on range of z^\perp . We first establish that vqv^* and m are partially orthogonal for every $m \in M$.

Let $m \in M$. As M was defined to be a partially orthogonal set of projections and $q \in M$, we know that q and m are partially orthogonal, and thus that $z^\perp q$ and $z^\perp m$ are partially orthogonal. However, as $z^\perp v = z^\perp$, we may express this as: $z^\perp(vqv^*)$

and $z^\perp m$ are partially orthogonal. Now, on the range of z , we have that

$$z(vqv^*) = z(uqu^*) \leq zs^R \quad \text{and} \quad zm \leq zs,$$

implying that $z(vqv^*)$ and zm are orthogonal, hence partially orthogonal. Putting both parts together, we have that $z^\perp vqv^*$ and $z^\perp m$ are partially orthogonal and that $z(vqv^*)$ and zm are partially orthogonal. We may thus apply Lemma 7.4.4 and conclude that vqv^* and m are partially orthogonal as desired.

Having established that vqv^* is partially orthogonal to all the projections in M , it follows by maximality of M that $vqv^* \in M$. Hence,

$$zvqv^* \leq vqv^* \leq \sup M = s.$$

Yet, by construction,

$$zvqv^* = zuqu^* \leq zs^R \leq s^R,$$

and so $zvqv^*$ must be orthogonal to s . Being both contained within and orthogonal to s , $zvqv^*$ must vanish. Therefore, the unitarily equivalent projection zq must also vanish. Now, $z \leq C(q)$ and $zq = 0$ forces z to be zero, for otherwise $C(q) - z$ would be both a central projection covering q and also strictly smaller than the central carrier of q . We may finally conclude that $s^R < uqu^*$. \square

7.5 Main theorem

Theorem 7.3.8, and thus our principal result, Theorem 7.3.6, will follow as an immediate corollary of:

Theorem 7.5.1. *In a von Neumann algebra \mathcal{A} , any invariant family of projections Φ arises from a central projection, i.e. Φ is equal to Φ_z for the central projection $z = \Phi(\mathcal{Z}(\mathcal{A}))$.*

Proof. Let Φ be an invariant family of projections. Suppose W is a commutative sub-von-Neumann-algebra of \mathcal{A} which contains the centre $\mathcal{Z}(\mathcal{A})$, and let q be the projection $\Phi(W)$. We claim that q is, in fact, equal to its own central carrier $C(q)$ and thus central. As $q \leq C(q)$ is true by definition, we must show that $q \geq C(q)$.

We start by applying Lemma 7.4.5 to q . Let M denote the resulting commuting set of projections in the unitary orbit of q , s denote the supremum of the projections in M , and s^R denote $C(q) - s = C(s) - s$. From the lemma, we know that $s^R \prec_u q$, i.e. there exists a unitary u such that $s^R < uqu^*$.

First note that, since $V_q \subset W$ and $q \in V_q$, by Lemma 7.4.1-iii, we have that $\Phi(V_q) = q$. Then, by unitary invariance of the family of projections, for every $m \in M$ we have that $\Phi(V_m) = m$. Hence, we can apply Lemma 7.4.2 to conclude that $\Phi(V_s) \geq s$. We also conclude, again by unitary invariance of Φ , that $\Phi(V_{uqu^*}) = uqu^* \geq s^R$.

Now, note that uqu^* and s^R commute and that $V_s = V_{s^R}$. So there is a commutative sub-von-Neumann-algebra $V_{s, uqu^*} \supseteq V_s, V_{uqu^*}$. By Lemma 7.4.1-i and the two conclusions of the preceding paragraph, we then have

$$\Phi(V_{s, uqu^*}) \geq \Phi(V_s) \vee \Phi(V_{uqu^*}) \geq s \vee s^R = C(q).$$

But, since $C(q) \in V_{uqu^*}$ by virtue of being contained in the centre, we can apply

Lemma 7.4.1-ii to find that $\Phi(V_{uqu^*}) \geq C(q)$. Finally, by unitary invariance,

$$q = \Phi(V_q) \geq u^*C(q)u = C(q) ,$$

concluding the proof that q is central.

We have shown that the projection $\Phi(W)$ is central for every commutative sub-von-Neumann-algebra W containing the centre $\mathcal{Z}(\mathcal{A})$. By Lemma 7.4.1-iii, this means that $\Phi(W)$ is equal to $\Phi(\mathcal{Z}(\mathcal{A}))$, the projection chosen at the centre, for all such W . In turn, this determines the image of Φ on all commutative sub-von-Neumann-algebras W' as

$$\Phi(W') = \sup \{ p \text{ is a projection in } W' \mid p \leq \Phi(V_{W' \cup \mathcal{Z}(\mathcal{A})}) = \Phi(\mathcal{Z}(\mathcal{A})) \} ,$$

and we find that Φ must be equal to $\Phi_{\Phi(\mathcal{Z}(\mathcal{A}))}$. □

7.6 Conclusions

Conjecture 7.3.5 is essentially the guess that the translation of the notion of closed set by the extension method matches up with the algebraic concept one would expect: closed, two-sided ideal. It would also recover the hull-kernel topology on the primitive ideal space of a C^* -algebra \mathcal{A} as a limit of the topologies; topologies of quotient spaces of the noncommutative space underlying \mathcal{A} .

We have established the von Neumann algebraic analogue of Conjecture 7.3.5. As a consequence, the original C^* -algebraic conjecture holds for all finite-dimensional C^* -algebras. The question of whether it holds for all C^* -algebras remains open. We conclude by indicating some ideas for future work that may lead to progress on this question.

One possible tack would be to enlarge the class of C^* -algebras for which the conjecture holds. An immediate suggestion would be the class of AF -algebras which arise as limits of finite-dimensional C^* -algebras [13]; it would follow immediately from a proof that \tilde{T} preserves limits.

Another possibility would be to prove the whole conjecture directly by using the proof of the von Neumann algebraic version as a guide. Indeed, one might still be able to reduce the question to one about projections by working in the enveloping von Neumann algebra \mathcal{A}^{**} of a C^* -algebra \mathcal{A} . In this setting, the total ideals of a C^* -algebra \mathcal{A} correspond to certain total ideals of the enveloping algebra \mathcal{A}^{**} [5]: those which correspond to open central projections. In essence, one would have to prove the appropriate analogue of Theorem 7.3.8 in order to find a correspondence between central open projections of \mathcal{A}^{**} and certain families of open projections which obey a restricted form of unitary invariance.

Chapter 8

Conclusions

With this dissertation, we have identified and explored a connection between the phenomenon of contextuality in quantum mechanics and the noncommutative geometry of C^* -algebras. We argued that, by taking state-observable duality seriously, a notion of *quantum state space* for a quantum system could be defined formally as being the geometric space such that the system's algebra of observables, which encapsulates all operational data, is manifested as the algebra of functions from the state space to outcomes. Since quantum systems are described by C^* -algebras of observables, construction of this quantum state space is tantamount to finding the geometric dual, in the categorical sense, to noncommutative C^* -algebras; to generalizing the Gel'fand spectrum functor from commutative to all C^* -algebras.

As the idea of representing observables as functions from states to outcomes (rather than to distributions on outcomes) is suggestive of determinism it is perhaps no surprise, in retrospect, that the phenomenon of contextuality plays a key role. Indeed, the idea of Isham and Butterfield's formulation in terms of context-indexed state spaces of the Kochen-Specker theorem provided a starting point for defining quantum state space. A hint which is sorely needed, as it is, of course, impossible to define a noncommutative spectrum in the most naive way: as a set of points with additional structure. The suggestion is then to consider as a possible notion of noncommutative spectrum not one topological space but a diagram of all the topological spaces which arise as quotient spaces of the hypothetical noncommutative space underlying a C^* -algebra.

Given this starting point, the field of noncommutative geometry gave a target. It provided generalizations of topological tools to the noncommutative setting which should, in principle, be straightforwardly reconstructible in a directly geometric way. A method for such reconstructions was proposed via an extension process which can be seen as decomposing a noncommutative space into its quotient spaces, retaining those which are genuine topological spaces, applying the topological func-

tor to each one, and pasting together the result.

It is our hope that the results of the last two chapters provide convincing evidence that the spatial diagrams associated to C^* -algebras in Chapter 4 are the closest topological approximations to noncommutative spaces. This immediately raises the question: *why* is the appropriate class of morphisms to consider those which arise from restricting inner automorphisms? Physically, these symmetries correspond to certain evolutions in time but how are they to be interpreted in terms of noncommutative topology?

More convincing evidence could be provided should some of the open problems to do with dropping the stability or finitary conditions in the K -theory result or the original C^* -algebraic conjecture about reconstructing the hull-kernel topology be resolved. Further examples of topological concepts which can be extended to their noncommutative analogues or useful applications of novel extensions would also strengthen the case.

To establish a concrete duality, it would be important to characterize which diagrams of spaces arise as spatial diagrams of a noncommutative algebra. For this perspective to be helpful for computations, some notion of a sub-spatial diagram ‘cover’ is needed. A useful analogy might be that of manifolds which come equipped with a maximal atlas. In practice, only a few charts are used for computations. Understanding how to recover a noncommutative algebra of functions from a spatial diagram is another key step.

The spatial diagram might also be explicitly calculated for some special examples. One promising possibility is the canonical commutation relations algebra which is closely connected with quantization and thus physically very significant. In this case, the Krichever-Mulase classification of certain commutative subalgebras of $\mathbb{C}[[x]][\partial]$ [54] provides a potentially highly useful roadmap. Another possibly tractable class of algebras for computations are those which arise as crossed prod-

uct algebras, wherein a group action on a C^* -algebra is embedded in a larger C^* -algebra such that the action is realized as a group of inner automorphisms. This class includes within it the important example of noncommutative tori (the computation of whose K -theory was considered a very difficult problem [59]).

The idea of looking at commutative quotients is a very general one and could perhaps be applied to analyze other sorts of noncommutative algebras other than C^* -algebras. The ideas outlined above might be applied to any duality involving a category of geometric objects and a category of commutative algebras.

The study of contextuality has extended beyond its original setting of quantum theory. These more general settings include test spaces or generalized probabilistic theories. This perspective of extending concepts defined for classical spaces to noncommutative spaces could provide a guide to finding appropriate analogues for concepts in classical probability theories in other general settings which exhibit contextuality.

Bibliography

- [1] Samson Abramsky and Adam Brandenburger. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13(11):113036, 2011.
- [2] J Frank Adams. Vector fields on spheres. *Annals of Mathematics*, pages 603–632, 1962.
- [3] Charles A. Akemann. Left ideal structure of C^* -algebras. *Journal of Functional Analysis*, 6(2):305–317, 1970.
- [4] Erik M. Alfsen. On the Dirichlet problem of the Choquet boundary. *Acta Mathematica*, 120(1):149–159, 1968.
- [5] Erik M. Alfsen and Frederic W. Shultz. *State Spaces of Operator Algebras: Basic theory, Orientations, and C^* -products*. Mathematics: theory & applications. Birkhäuser, 2001.
- [6] Erik M. Alfsen and Frederick W. Shultz. *Geometry of State Spaces of Operator Algebras*. Mathematics (Boston, Mass.). Birkhäuser Boston, 2003.
- [7] Michael F. Atiyah and Donald W. Anderson. *K-theory*. WA Benjamin New York, 1967.
- [8] Michael F. Atiyah and Raoul Bott. On the periodicity theorem for complex vector bundles. *Acta Mathematica*, 112(1):229–247, 1964.

- [9] Michael Barr. The Chu construction. *Theory and Applications of Categories*, 2(2):17–35, 1996.
- [10] John S. Bell. On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics*, 38(3):447, 1966.
- [11] Klaus Bichteler. A generalization to the non-separable case of Takesaki’s duality theorem for C^* -algebras. *Inventiones Mathematicae*, 9(1):89–98, 1969.
- [12] Bruce Blackadar. *Operator Algebras - Theory of C^* -Algebras and von Neumann Algebras*. Number 122-III in Encyclopaedia of Mathematical Sciences: Operator Algebras and Non-Commutative Geometry (Joachim Cuntz and Vaughan F. R. Jones eds.). Springer, 2006.
- [13] Ola Bratteli. Inductive limits of finite dimensional C^* -algebras. *Transactions of the American Mathematical Society*, 171:195–234, 1972.
- [14] Martijn Caspers, Chris Heunen, N. P. Landsman, and Bas Spitters. Intuitionistic quantum logic of an n -level system. *Foundations of Physics*, 39:731–759, 2009.
- [15] Alain Connes. *Noncommutative Geometry*. Springer, 1992.
- [16] John Dauns and Karl Heinrich Hofmann. *Representation of Rings by Sections*, volume 83. American Mathematical Society, 1968.
- [17] Hans F. de Groote. Observables IV: The presheaf perspective. *arXiv preprint arXiv:0708.0677*, 2007.
- [18] Andreas Döring. Kochen-Specker theorem for von Neumann algebras. *International Journal of Theoretical Physics*, 44(2):139–160, 2005.
- [19] Andreas Döring and Chris J. Isham. A topos foundation for theories of physics: I. Formal languages for physics. *Journal of Mathematical Physics*, 49(5):053515, 2008.

- [20] Andreas Döring and Chris J. Isham. A topos foundation for theories of physics: II. Daseinisation and the liberation of quantum theory. *Journal of Mathematical Physics*, 49(5):053516, 2008.
- [21] Andreas Döring and Chris J. Isham. A topos foundation for theories of physics: III. Quantum theory and the representation of physical quantities with arrows $\check{A} : \underline{\Sigma} \rightarrow \mathcal{P}\mathbb{R}$. *Journal of Mathematical Physics*, 49(5):053517, 2008.
- [22] Andreas Döring and Chris J. Isham. A topos foundation for theories of physics: IV. Categories of systems. *Journal of Mathematical Physics*, 49(5):053518, 2008.
- [23] Samuel Eilenberg and Norman E Steenrod. Axiomatic approach to homology theory. *Proceedings of the National Academy of Sciences of the United States of America*, 31(4):117, 1945.
- [24] George A. Elliott. The classification problem for amenable C^* -algebras. In *Proceedings of the International Congress of Mathematicians*, pages 922–932. Springer, 1995.
- [25] J.M.G. Fell. The structure of algebras of operator fields. *Acta Mathematica*, 106(3):233–280, 1961.
- [26] Peter A. Fillmore. *A User’s Guide to Operator Algebras*. Wiley New York, 1996.
- [27] Henrik Forssell. *First-order logical duality*. PhD thesis, Carnegie Mellon University, 2008.
- [28] Ichiro Fujimoto. A Gelfand–Naimark theorem for C^* -algebras. *Pacific Journal of Mathematics*, 184(1):95–119, 1998.
- [29] Israïl M. Gel’fand and Mark A. Naïmark. On the imbedding of normed rings into the ring of operators in Hilbert space. *Rec. Math. [Mat. Sbornik] N.S.*, 12(54)(2):197–213, 1943.

- [30] Robin Giles and Hans Kummer. A non-commutative generalization of topology. *Indiana University Mathematics Journal*, 21(1):91–102, 1972.
- [31] Andrew M. Gleason. Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6(6):885–893, 1957.
- [32] Alexandre Grothendieck. *Classes de faisceaux et théoreme de Riemann-Roch*. 1957.
- [33] Jan Hamhalter. *Quantum Measure Theory*, volume 134. Springer, 2003.
- [34] John Hamilton, Chris J. Isham, and Jeremy Butterfield. A topos perspective on the Kochen-Specker theorem: III. Von Neumann algebras as the base category. *International Journal of Theoretical Physics*, 39(6):1413–1436, 2000.
- [35] Robin Hartshorne. *Algebraic Geometry*. Number 52. Springer, 1977.
- [36] Carsten Held. The Kochen-Specker theorem. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Spring 2013 edition, 2013.
- [37] Chris Heunen, N. P. Landsman, and Bas Spitters. A topos for algebraic quantum theory. *Communications in Mathematical Physics*, 291:63–110, 2009.
- [38] Chris Heunen, N. P. Landsman, and Bas Spitters. Bohrification of operator algebras and quantum logic. *Synthese*, 2010.
- [39] Chris Heunen, N. P. Landsman, and Bas Spitters. Bohrification. In Hans Halvorson, editor, *Deep Beauty*. Cambridge University Press, 2011.
- [40] Chris Heunen, N. P. Landsman, Bas Spitters, and Sander Wolters. The Gelfand spectrum of a noncommutative C^* -algebra: a topos-theoretic approach. *Journal of the Australian Mathematical Society*, 90(01):39–52, 2011.
- [41] Chris J. Isham. *Lectures on Quantum Theory - Mathematical and Structural Foundations*. Imperial College Press, 1995.

- [42] Peter T. Johnstone. *Stone Spaces*, volume 3 of *Cambridge studies in advanced mathematics*. Cambridge University Press, 1982.
- [43] Richard V. Kadison. *A Representation Theory for Commutative Topological Algebra*. Number 7. American Mathematical Society, 1951.
- [44] Richard V. Kadison and John R. Ringrose. *Fundamentals of the Theory of Operator Algebras*, volume I - Elementary Theory. Academic Press, 1983.
- [45] Masoud Khalkhali. Very basic noncommutative geometry. *ArXiv Mathematics e-prints*, August 2004.
- [46] Simon Kochen and Ernst P. Specker. The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17(1):59–87, 1967.
- [47] David Kruml, Joan Wick Pelletier, Pedro Resende, and Jiří Rosický. On quantales and spectra of C^* -algebras. *Applied Categorical Structures*, 11(6):543–560, 2003.
- [48] Pawel Kruszynski and Stanislaw L. Woronowicz. A non-commutative Gelfand-Naimark theorem. *Journal of Operator Theory*, 8:361–389, 1982.
- [49] Yeong-Cherng Liang, Robert W. Spekkens, and Howard M. Wiseman. Specker’s parable of the overprotective seer: A road to contextuality, non-locality and complementarity. *Physics Reports*, 506(1):1–39, 2011.
- [50] Saunders Mac Lane. *Categories for the working mathematician*. Number 5 in Graduate Texts in Mathematics. Springer-Verlag, 1971.
- [51] George W. Mackey. Quantum mechanics and Hilbert space. *The American Mathematical Monthly*, 64(8):pp. 45–57, 1957.
- [52] George W. Mackey. *Mathematical Foundations of Quantum Mechanics*. W. A. Benjamin, 1963.

- [53] Jerrold E. Marsden and Tudor S. Ratiu. *Introduction to Mechanics and Symmetry: A Basic Exposition of Classical Mechanical Systems*, volume 17. Springer, 1999.
- [54] Motohico Mulase. Category of vector bundles on algebraic curves and infinite dimensional Grassmannians. *International Journal of Mathematics*, 1(03):293–342, 1990.
- [55] Christopher J. Mulvey. A non-commutative Gel’fand-Naimark theorem. *Mathematical Proceedings of the Cambridge Philosophical Society*, 88:425–428, 11 1980.
- [56] Gert K. Pedersen. SAW^* -algebras and corona C^* -algebras, contributions to non-commutative topology. *Journal of Operator Theory*, 15:15–32, 1986.
- [57] Pedro Resende. Étale groupoids and their quantales. *Advances in Mathematics*, 208(1):147–209, 2007.
- [58] Manuel L. Reyes. Obstructing extensions of the functor Spec to noncommutative rings. *Israel Journal of Mathematics*, 192(2):667–698, 2012.
- [59] Marc Rieffel. C^* -algebras associated with irrational rotations. *Pacific Journal of Mathematics*, 93(2):415–429, 1981.
- [60] Mikael Rørdam, Flemming Larsen, and Niels Laustsen. *An Introduction to K-theory for C^* -algebras*, volume 49. Cambridge University Press, 2000.
- [61] Terry Rudolph. Ontological Models for Quantum Mechanics and the Kochen-Specker theorem. *eprint arXiv:quant-ph/0608120*, August 2006.
- [62] Mike Shulman and Urs Schreiber. nLab: slice 2-category, 2013.
- [63] Frederic W. Shultz. Pure states as a dual object for C^* -algebras. *Communications in Mathematical Physics*, 82(4):497–509, 1982.

- [64] Robert W. Spekkens. Contextuality for preparations, transformations, and unsharp measurements. *Physical Review A*, 71:052108, May 2005.
- [65] Marshall Stone. The theory of representations for Boolean algebras. *Transactions of the American Mathematical Society*, 40:37–111, 1936.
- [66] Franco Strocchi. *An Introduction to the Mathematical Structure of Quantum Mechanics: A Short Course for Mathematicians*. World Scientific, 2008.
- [67] Richard G. Swan. Vector bundles and projective modules. *Transactions of the American Mathematical Society*, pages 264–277, 1962.
- [68] Masamichi Takesaki. A duality in the representation theory of C^* -algebras. *Annals of Mathematics*, pages 370–382, 1967.
- [69] Tammo tom Dieck. *Algebraic Topology*. European Mathematical Society, 2008.
- [70] John von Neumann. *Mathematische Grundlagen der Quantenmechanik*. Springer, 1932.
- [71] Niels Erik Wegge-Olsen. *K-theory and C^* -algebras: A Friendly Approach*. Oxford Science Publications. Oxford University Press, 1993.