

Projective Limits of State Spaces: Quantum Field Theory Without a Vacuum

Suzanne Lanéry*

Centro de Ciencias Matemáticas, Universidad Nacional Autónoma de México, Apartado Postal 61-3, 58089 Morelia, Michoacán, México

Received 27 December 2017, Accepted 16 March 2018, Published 20 April 2018

Abstract: Instead of formulating the states of a Quantum Field Theory (QFT) as density matrices over a single large Hilbert space, it has been proposed by Kijowski [20] to construct them as consistent families of *partial* density matrices, the latter being defined over small 'building block' Hilbert spaces. In this picture, each small Hilbert space can be physically interpreted as extracting from the full theory specific degrees of freedom. This allows to reduce the quantization of a classical field theory to the quantization of *finite-dimensional* sub-systems, thus sidestepping some of the common *ambiguities* (specifically, the issues revolving around the choice of a 'vacuum state'), while obtaining robust and well-controlled quantum states spaces. The present letter provides a self-contained introduction to this formalism, detailing its motivations as well as its relations to other approaches to QFT (such as conventional Fock-like Hilbert spaces, path-integral quantization, and the algebraic formulation). At the same time, it can serve as a reading guide to the series of more in-depth articles [27, 28, 29, 30].

© Electronic Journal of Theoretical Physics. All rights reserved.

Keywords: Quantum Field Theory; Vacuum States; Inequivalent Representations; Geometric Quantization; Projective Limits; Algebras of Observables

PACS (2010): 02.40.Yy; 03.50.-z; 03.70.+k; 04.62.+v; 04.60.Ds

1. Motivation: Quantization Ambiguities in Quantum Field Theory

Many choices have to be made in the quantization of a classical theory. Assuming one is following the canonical quantization path (see section 5. for further discussion of the relevance for path-integral approaches of the issues discussed here), the first step is to choose a complete set of basic variables for the theory. Heuristically, these are the vari-

* Email: slanery@matmor.unam.mx

ables for which the semi-classical limit will work best, hence their choice should ideally reflect the observables against which the classical theory of interest has been best tested and confirmed.

The next step is to find a representation of these basic variables as operators on a suitable Hilbert space \mathcal{H} , namely a mapping $f \mapsto \hat{f}$ such that

$$[\hat{f}, \hat{g}] = i \widehat{\{f, g\}} \quad (1)$$

(where $[\cdot, \cdot]$ denotes the commutator of operators, while $\{\cdot, \cdot\}$ denotes the Poisson brackets of classical observables). At this point, quantum field theory (in a broad sense, namely quantum theories meant to encompass *infinitely* many degrees of freedom) differs crucially from quantum mechanics (dealing with the quantum counterparts of classical systems that have *finitely* many degrees of freedom). The tools from geometric quantization [45] (that we will discuss further in subsection 2.2) provide a clear and detailed understanding of the canonical quantization of finite dimensional systems, including a parametrization of available choices (aka. quantization ambiguities). In some cases, it may even turn out that there is no choice at all, because the Poisson-algebra of interest admits only one suitable representation: this is for example the content of the Stone-von-Neumann theorem [41, 42, 39] in the case of linear systems.

By contrast, the representation theory for infinite dimensional system tends to be very involved. Even in the simplest case of a free scalar field on Minkowski spacetime, it is known that there exist infinitely many *inequivalent* representations, and although it has been possible, in this very special case, to fully classify them [15], this classification is so complex that it gives little insight on how to choose one. As a way out, a pragmatic way of selecting a good representation among these too numerous options is to single out a distinguished quantum state, the *vacuum*: it is indeed possible, via the so-called GNS construction [11, 38] to 'seed' a full representation \mathcal{H}_Ω from a single state Ω (to specify the latter, even before we are equipped with a Hilbert space, we can give the corresponding expectation values of all products of the basic variables, aka. the n -point functions, see [17, part III, def. 2.2.8]). This approach has established itself as the standard way to think about quantum field theory, at least in the context of Minkowski spacetime, where the vacuum may be selected by requiring it to be invariant under all spacetime symmetries (ie. under the Poincaré group).

However, one should keep in mind that the initial choice of vacuum is deeply imprinted in the thus obtained representation. The only quantum states that can be written as (pure or statistical) states on \mathcal{H}_Ω are those that barely differ from the vacuum: at most discrete quantum excitations on top of the state Ω are allowed. The set of all states living on the representation \mathcal{H}_Ω is referred to as the vacuum *sector*, in acknowledgment of the fact that there are many more quantum states beyond it (falling out of it because they lie too far away from the chosen vacuum), among whose some may actually be interesting for specific purposes [17, part V]. An implication of the relative smallness of the vacuum sector is that the vacuum state need to be closely tailored to the *dynamics*: otherwise, the time evolution would immediately kick the states out of \mathcal{H}_Ω (a precise statement of this

heuristic expectation is given, for Poincaré-invariant QFTs, by the Haag no-go theorem, [16]).

A radical alternative, prompted by the lack of a natural vacuum in the case of quantum field theory on *curved* spacetime, is to use as state space the *whole* set of possible quantum states over the chosen basic observables (each such state being specified, as explained above, by the expectations values it prescribes for all products of observables). This approach can be followed in the context of Algebraic Quantum Field Theory (AQFT, [17, 18]): by shifting the focus from a particle picture to the *local* and *causal* structure of the quantum theory, AQFT provides tools to discuss the properties of quantum fields in the absence of an underlying Hilbert space. The aim of the present letter is to argue that a projective definition of quantum field theory, as was introduced by Jerzy Kijowski [20] and further developed by Andrzej Okołów [32, 34, 33], can provide a middle way between the conventional vacuum-based approach and the full algebraic one, retaining a *constructive* description of the quantum state space (subsections 3.2 and 3.3) while keeping enough *flexibility* to accommodate a wide class of quantum states (subsection 3.1) and to decouple the subsequent implementation of the dynamics from the initial building of the state space (section 4.).

The work summarized in the following sections (and developed in details in [27, 28, 29, 30]) was notably motivated by the specific difficulties encountered when one tries to formulate *background independent* quantum field theories, rather than theories on a (possibly curved) background spacetime (eg. to quantize general relativity itself in a non-perturbative way [2, 40]). It turns out that for background independent gauge theories (at least those with *compact* gauge group), there does exist a *preferred* vacuum state, the Ashtekar-Lewandowski vacuum [3, 4], which is uniquely selected precisely by the requirement of background independence [31, 10]. Unfortunately, this vacuum has some unwanted properties. One of them is that it is an *eigenstate* of the variable conjugate to the gauge field, rather than a *coherent state* like the usual Fock vacuum. Since states in the vacuum sector cannot differ too much from the vacuum, this makes it difficult to find semi-classical states among them [12, 22]. Another problem is that the GNS representation built on this vacuum lives on a *non-separable* Hilbert space. This particular issue may or may not go away once we identify quantum states that only differ by a change of coordinates (depending on how precisely this identification is carried out, see [40, 8]) but in any cases it can lead to technical difficulties [5]. Paradoxically, non-separable Hilbert spaces seem too *small*, because their orthonormal basis need uncountably many basis vectors, making it tempting to consider uncountable linear combinations while only countable ones are allowed: in other words, it is in this case even more likely that physically interesting states will lie out of the vacuum sector. In an effort to overcome these difficulties, the projective quantization techniques that we will review in the next section have been applied to this kind of theories, resulting in a quantum state space that may have applications to the study of the semi-classical and cosmological sectors of quantum gravity [26].




Fig. 1 Three-spaces consistency for projective systems (left side), reformulated in terms of factorizations (right side)

2. Systematic Quantization of Infinite-dimensional Systems

2.1 Building an Infinite-dimensional Theory from a Collection of Partial Descriptions

The key observation underlying Kijowski's projective formalism [20, 33] is that a given experiment can only measure finitely many observables. Thus, we never need to consider at once the full, *infinite*-dimensional phase space \mathcal{M}_{∞} of a field theory: it is sufficient to work in a small, partial phase space \mathcal{M}_{η} that extracts from \mathcal{M}_{∞} just the degrees of freedom (dof.²) relevant for the experiment at hand (throughout the present letter, the symbol η will be used to denote a selection of finitely many dof. out of the full theory, and we will call η a *label*).

In order to use such a collection of finite-dimensional partial phase spaces $(\mathcal{M}_{\eta})_{\eta}$ to completely specify a field theory, we need to ensure that the different partial theories are consistent with each other [27, subsection 2.1]:

1. first, we need a way to express the relations between the dof. in different labels. We will write $\eta \preceq \eta'$ if all dof. contained in η are also contained in η' (we will also say that η is *coarser* as η' , or that η' is *finer* as η). This means that any observable f_{η} on \mathcal{M}_{η} corresponds to an observable $f_{\eta'}$ on $\mathcal{M}_{\eta'}$, and, by duality³, that there exists a projection $\pi_{\eta' \rightarrow \eta}$ from $\mathcal{M}_{\eta'}$ to \mathcal{M}_{η} such that

$$f_{\eta'} = f_{\eta} \circ \pi_{\eta' \rightarrow \eta}$$

2. the predictions for a given experiment, as calculated in a partial theory η , should be independent of the choice of η (provided η is fine enough to hold all relevant dof.). Thus, in particular, the Poisson brackets between two observables f_{η} and g_{η} on \mathcal{M}_{η}

² By a dof. we mean a *pair* of conjugate variables.

³ To see that $\pi_{\eta' \rightarrow \eta}$ is uniquely specified once we know the mapping $f_{\eta} \mapsto f_{\eta'}$ between observables, one can consider a complete set of observables (aka. coordinates) on \mathcal{M}_{η} .

should agree with the Poisson brackets between the corresponding observables on a finer $\mathcal{M}_{\eta'}$. Expressed in terms of the just introduced projection $\pi_{\eta' \rightarrow \eta}$, this reads

$$\{f_{\eta} \circ \pi_{\eta' \rightarrow \eta}, g_{\eta} \circ \pi_{\eta' \rightarrow \eta}\} = \{f_{\eta'}, g_{\eta'}\} \circ \pi_{\eta' \rightarrow \eta} \quad (2)$$

3. it should be possible to consider composite experiments made of two (or more) sub-experiments, each of which can be described within a different partial theory⁴. In other words, for any η, η' , there should exist η'' such that $\eta, \eta' \preceq \eta''$. This property is called *directedness* of the set of labels.
4. the relation between any two partial theories should be unambiguous. Thanks to the just mentioned directedness property, this can be ensured simply by requiring that the projections defined among three increasingly refined partial theories match as shown on the left part of fig. 1.

In mathematical terms, this list of requirements can be summarized by saying that the collection $(\mathcal{M}_{\eta})_{\eta}$ forms a projective (aka. inverse) system and it ensures that \mathcal{M}_{∞} can be reconstructed from $(\mathcal{M}_{\eta})_{\eta}$ (more precisely, a space \mathcal{M}_{lim} can be constructed as the so-called projective limit of this system, that will, in general, be a distributional extension of \mathcal{M}_{∞} , see [27, def. 2.6 and prop. 2.7]).

A key point of the construction is that each label η corresponds to a selection of conjugate position *and* momentum variables. This ensures that the projection map $\pi_{\eta' \rightarrow \eta}$ between the phase spaces $\mathcal{M}_{\eta'}$ and \mathcal{M}_{η} is unambiguous: as stressed above, it is completely determined by matching the physical interpretation of the observables in η vs. η' . By contrast, if the partial theories were labeled by selections of configuration variables only (as is usual eg. when studying coarse graining in the path-integral formalism), we would only be provided with projections between the configuration spaces (since phase spaces can be thought of as *cotangent* bundles, and forms are naturally pull-backed, rather than push-forwarded, there is no canonical prescription to lift a projection between configuration spaces to a projection between the associated phase spaces). Furthermore, our labels do *not* have to be assembled from pre-assigned pairs of mutually independent, canonically conjugate variables (aka. modes). In particular, different labels can consist of the *same* selection of configuration variables, but paired with *different* conjugate impulsions⁵: the corresponding partial theories are then interpreted as extracting from the full, continuum field theory information about different (partially overlapping) dof. The directedness requirement (which is much weaker than the requirement of a preferred mode decomposition, and is indeed fulfilled in a range of situations where the latter would not

⁴ We are discussing the classical theory here. The quantum theory is more subtle, since one could argue that, due to the principle of complementarity, some sub-experiments may be mutually excluding. However, the case for the directedness of the label set can still be made, see [28, section 1].

⁵ This is manifest eg. in the treatment of gauge theories in [26]: selections of configuration variables are represented by graphs (generalized lattices) and these need to be decorated with dual surfaces to fully specify the impulsions (aka. electric fluxes) captured by a given partial theory (see the discussion at the beginning of [26, section II]).




Fig. 2 Factorization from a Poisson-brackets-preserving projection between phase spaces

be an option; see [26, 25]) guarantees that these partial informations can be assembled into a consistent picture.

A useful consequence of eq. (2) is that it ensures that there exists a *preferred factorization* of $\mathcal{M}_{\eta'}$ as $\mathcal{M}_{\eta' \rightarrow \eta} \times \mathcal{M}_{\eta}$ (at least locally: there may be some topological obstructions preventing this factorization to hold globally). The way this factorization is obtained is as follow⁶ (see [33, section 3.4] and [27, prop. 2.10]). Any observable f_{η} on \mathcal{M}_{η} generates a Hamiltonian flow Φ_{η}^t on \mathcal{M}_{η} , and $f_{\eta'}$ generates a corresponding Hamiltonian flow $\Phi_{\eta'}^t$ on $\mathcal{M}_{\eta'}$. All the flows on $\mathcal{M}_{\eta'}$ generated in this manner from all possible observables on \mathcal{M}_{η} together span joint orbits in $\mathcal{M}_{\eta'}$ (represented as dashed lines on fig. 2), and eq. (2) ensures that any such orbit can (again, locally) be identified with \mathcal{M}_{η} . All what is left to do is then to define $\mathcal{M}_{\eta' \rightarrow \eta}$ as the corresponding quotient, ie. the set of all such orbits (see fig. 2). $\mathcal{M}_{\eta' \rightarrow \eta}$ can thus be interpreted as holding dof. that are complementary to the ones retained by \mathcal{M}_{η} (since, by construction the dof. in $\mathcal{M}_{\eta' \rightarrow \eta}$ Poisson-commute with the ones in \mathcal{M}_{η}).

If we assume, for simplicity, that, for any $\eta \preceq \eta'$, the factorization $\mathcal{M}_{\eta'} \approx \mathcal{M}_{\eta' \rightarrow \eta} \times \mathcal{M}_{\eta}$ holds *globally*, we can translate the consistency condition illustrated on the left part of fig. 1 into a condition written in terms of the factorizations [27, def. 2.11], as shown on the right part of fig. 1. The meaning of this condition is that, when extracting from η'' the dof. corresponding to a much coarser partial theory η , the complementary dof. that we discard should be the same whether we go from η'' to η in one step or in two, hence there should be a natural identification $\mathcal{M}_{\eta'' \rightarrow \eta} \approx \mathcal{M}_{\eta' \rightarrow \eta} \times \mathcal{M}_{\eta'' \rightarrow \eta'}$.

⁶ Again, it is crucial that the η 's are labeling *phase spaces*: the Poisson-bracket structure is what allows to identify the complementary observables. By contrast, if we were only provided with a projection between configuration spaces, there would be no way to single out a preferred factorization of the larger space.

2.2 Leveraging the Well-developed Procedures for Quantizing Finite-dimensional Systems

Once a collection of *classical* partial theories has been put in the just mentioned factorized form, constructing a corresponding collection of *quantum* partial theories is straightforward [28]. All we have to do is to quantize each \mathcal{M}_η into a corresponding Hilbert space \mathcal{H}_η , and, since \mathcal{M}_η is *finite*-dimensional, we have at our disposal the full range of techniques known for quantizing systems with finitely many dof. From the usual rules of quantum mechanics we expect that, for any $\eta \preceq \eta'$, the factorization of phase spaces $\mathcal{M}_{\eta'} \approx \mathcal{M}_{\eta' \rightarrow \eta} \times \mathcal{M}_\eta$ will give rise to a tensor-product factorization of Hilbert spaces $\mathcal{H}_{\eta'} \approx \mathcal{H}_{\eta' \rightarrow \eta} \otimes \mathcal{H}_\eta$.

In this way, we can obtain a collection of Hilbert spaces, with relations between them expressed in terms of tensor-products. But where is the promised quantum state space for the full field theory? The crucial realization, that goes back to Kijowski [20], is that we can now construct quantum states as collections $(\rho_\eta)_\eta$, where each ρ_η is a density matrix (aka. mixed or statistical state) on \mathcal{H}_η . In order for such a collection to describe a quantum state for the full theory, we need to ensure that the partial states ρ_η are consistent with each other. Namely, whenever $\eta \preceq \eta'$, the states ρ_η and $\rho_{\eta'}$ should lead to the same measurement probabilities for all dof. retained by the coarser label η . Using well-known tools from statistical quantum mechanics, this means that we need to enforce

$$\rho_\eta = \text{Tr}_{\eta' \rightarrow \eta} \rho_{\eta'} \quad (3)$$

where $\text{Tr}_{\eta' \rightarrow \eta}$ denotes the *partial trace* on the tensor product factor $\mathcal{H}_{\eta' \rightarrow \eta}$. Indeed, the partial trace satisfies the important property

$$\text{Tr} \rho_\eta \hat{f}_\eta = \text{Tr} (\text{Tr}_{\eta' \rightarrow \eta} \rho_{\eta'}) \hat{f}_\eta = \text{Tr} \rho_{\eta'} (1_{\eta' \rightarrow \eta} \otimes \hat{f}_\eta) = \text{Tr} \rho_{\eta'} \hat{f}_{\eta'} \quad (4)$$

In order to be able to consistently impose eq. (3), the way to reconstruct a coarser partial quantum state from a finer one should be unambiguous, in other words we should have

$$\text{Tr}_{\eta'' \rightarrow \eta} = \text{Tr}_{\eta' \rightarrow \eta} \circ \text{Tr}_{\eta'' \rightarrow \eta'}$$

for any three increasingly fine labels $\eta \preceq \eta' \preceq \eta''$. Fortunately, this is automatically guaranteed by the quantum analogue of the classical consistency condition from fig. 1 [28, defs. 2.1 and 2.2]. The quantum state space obtained in this way may be called *projective*, since it arises mathematically as a projective limit (like the space \mathcal{M}_{lim} mentioned in subsection 2.1).

Note that it is necessary in this construction to use density matrices rather than Hilbert space vectors (aka. pure states). Indeed, the partial trace is the right tool to use when one is interested in the restriction of a quantum state to specific dof. (thanks to the property stressed in eq. (4)) and it *cannot* be defined as a map between pure states, because the partial trace of a pure state is often a mixed state and, reciprocally, the partial trace of a mixed state can appear pure⁷. Another way of looking at this is

⁷ When $\rho_{\eta'}$ has the special form $\tilde{\rho} \otimes |\psi\rangle\langle\psi|$, any statistical superpositions in $\tilde{\rho}$ will be traced out.

that the distinction between pure states and mixed ones is always relative to a certain (sub)system, hence it is effectively useless in an infinite-dimensional theory: the only absolute notion of state purity would be the one defined with respect to the full theory, but such an absolute notion cannot have any experimental relevance, because it would require to measure quantum states against a *complete* set of commuting observables, that is, to perform *infinitely* many measurements.

In [28, section 3] and [25] more specific results have been obtained regarding the implementation of this general strategy. As is apparent from the discussion above, what we need is a *consistent* quantization scheme, namely a prescription to quantize each partial theory \mathcal{M}_η in such a way that the Cartesian product factorizations, holding on the classical side, are indeed lifted to tensor-product factorizations on the quantum side, and that the quantization of observables agree, ie. that we do have $\hat{f}_{\eta'} = 1_{\eta' \rightarrow \eta} \otimes \hat{f}_\eta$ as was assumed in eq. (4).

Such a scheme can be adapted from geometric quantization [45]. The latter formalizes the traditional recipes of canonical quantization into a systematic procedure, which can be understood starting from the following observation: A very simple way of quantizing a finite dimensional phase space \mathcal{M}_η would be to set $\mathcal{H}_\eta = L_2(\mathcal{M}_\eta)$ ⁸ and, for any observable f_η on \mathcal{M}_η

$$\hat{f}_\eta \psi = -i \{f_\eta, \psi\}.$$

Then eq. (1) is simply the Jacobi equation satisfied by the Poisson brackets. This naive ansatz can be improved by incorporating some multiplicative action of f_η in the definition of \hat{f}_η , yielding what is known as *pre-quantization* [45, section 8.2]. Since the pre-quantization primarily depends on the Poisson bracket structure of the classical theory, and this structure is preserved by the projections $\pi_{\eta' \rightarrow \eta}$ as required in subsection 2.1, it generically provides a consistent quantization scheme in the sense above [28, theorem 3.9 and prop. 3.11].

Pre-quantization is however not fully satisfactory (hence the 'pre' qualification) as it leads to Hilbert spaces which are simply too big. Physically admissible Hilbert spaces are extracted by restricting ourselves to a subset of the pre-quantized Hilbert space [45, chap. 9]: for example the usual *position* representation of quantum mechanics is recovered by considering only the pre-quantum states $\psi(x, p)$ which are constant in p , ie. only depend on x , while *holomorphic* quantization is obtained by keeping only the ψ which are holomorphic functions (with respect to a suitable complex structure on \mathcal{M}_η , eg. $z = x + ip$). This extra step of the geometric quantization program is called imposing a *polarization*. For the purpose of constructing projective quantum state spaces, the question is therefore whether polarizations can be imposed *consistently* on each pre-quantized \mathcal{H}_η .

Building on previous work by Okołów [33], this question was examined in two concrete cases (namely, position quantization when the configuration spaces for all partial theories

⁸ Note that \mathcal{M}_η , as a phase space, admits a natural measure: the Liouville measure, which is singled out (up to a multiplicative constant) and has the important property of being preserved under any Hamiltonian flow, which ensures, in particular, that \hat{f}_η will be self-adjoint whenever f_η real.

are Lie groups and the projective structure from subsection 2.1 is suitably compatible with the group structure, in [28, subsection 3.1], and holomorphic quantization when all projection maps $\pi_{\eta' \rightarrow \eta}$ are holomorphic, in [28, subsection 3.2]) and the answer was found to be positive. These results suggest that polarizations can be consistently imposed on the pre-quantized projective system whenever the classical precursors of these polarizations are compatible with the coarse-graining projections $\pi_{\eta' \rightarrow \eta}$. Furthermore, the nature of the compatibility required turns out to be rather weak, because it only has to be satisfied up to *unitary equivalence* of the quantizations. For example, if the small, finite-dimensional phase spaces \mathcal{M}_η are *linear* spaces, the Stone–von Neumann theorem mentioned at the beginning of section 1. ensures that all polarizations which respect the linear structure lead to unitarily equivalent representations \mathcal{H}_η . Leveraging these equivalences, a general quantization scheme for linear (aka. free) field theories was developed in [25], which is *explicitly independent* of any overall choice of polarization on the (infinite-dimensional) classical phase space: specifically, the label set it uses spans all possible choices of polarization at once, so that none has to be singled out, making it suitable, eg. to the study of quantum fields on curved backgrounds (while field theories defined on a flat background spacetimes admits a physically preferred polarization, this is no longer the case on curved backgrounds, especially non-static ones; see [18]).

Moreover, it may also be possible to dispense from the requirement that the classical factorizations $\mathcal{M}_{\eta'} \approx \mathcal{M}_{\eta' \rightarrow \eta} \times \mathcal{M}_\eta$ hold *globally* (see [28, section 4]).

3. Relations to Standard Constructions: Improved Universality and Richer State Space

3.1 Embedding All Vacuum Sectors at Once

Since the program exposed in the previous section is meant to improve on the vacuum-based quantization described in section 1., we need to understand how the quantum state spaces obtained by both methods relate to each other. Suppose that we are given a vacuum state Ω , which can be completely specified, as underlined in section 1., by the expectation values it prescribes for any product of observables. Then, we can define a partial vacuum state Ω_η , by simply forgetting about the expectation value of any product involving an observable that does not live on the partial theory labeled by η . Now, quantum theories describing finitely many dof. do not suffer from the pathologies affecting quantum field theory: the states that can be written as density matrices on the Hilbert space \mathcal{H}_η are likely to span a large part of the space of possible quantum states. Thus, it is not unreasonable to assume that each Ω_η corresponds to a density matrix ρ_η on \mathcal{H}_η , and, since the Ω_η all stem from the same state Ω of the full theory, they are guaranteed to be consistent, ensuring that eq. (3) holds. In other words, the vacuum Ω is likely to belong to the projective quantum state space constructed in subsection 2.2, and with him all quantum states that only differ marginally from Ω , ie. the whole vacuum sector around Ω [28, prop. 2.8]. This reasoning suggests that the vacuum sectors corresponding

to a large class of vacuums are included for free in the projective state space.

In [28, theorem 2.9], we examined in particular the case where the ρ_η obtained from Ω are actually *pure* states for all η . Note that in contrast to the discussion so far, this is a very special situation, in the sense that pure states are fairly ‘exceptional’ in the space of all quantum states [17, part III, theorem 2.2.18]; yet, it turns out to apply to the most interesting cases (such as the embedding of the usual Fock space, or of the Ashtekar-Lewandowski state space mentioned in section 1., into their projective counterparts, see [29, prop. 3.17] and [25, subsection 3.3], or [26, theorem 40], respectively). Under this additional assumption, it is possible to prove that the embedding of the corresponding vacuum sector into the projective state space is injective (aka. one-to-one) and its image can be precisely characterized.

A different construction that could claim to include at once many vacuum sectors are the so-called infinite tensor products (ITP, [43]). These can be defined as huge direct sums of GNS Hilbert spaces \mathcal{H}_Ω (recall section 1.) for Ω running through a suitable class of vacuums. In other words, the different vacuum sectors are simply put *side by side*. Besides resulting in a large, non-separable Hilbert space, a drawback of this construction is that it is fairly fragile: it requires a decomposition of the dof. into independent (commuting) subsets, with the slightest modification of this decomposition leading to *inequivalent* ITP Hilbert spaces. More seriously, it introduces spurious distinctions between physically indistinguishable states: since the different sectors appear as *superselection sectors* (ie. the observables are not sensible to quantum correlations *between* these sectors), quantum superpositions of states in different sectors are indistinguishable from statistical superpositions. By contrast, projective state spaces, while allowing for statistical superpositions of states in different vacuum sectors, do not produce such an unphysical redundancy. This is exhibited by the fact that we can map the states on an ITP to states in a corresponding projective state space, in such a way that ITP states are mapped to the *same* projective state whenever they are physically indistinguishable [28, theorem 2.11].

3.2 The Case for Countable Collections of Partial Theories

While projective state spaces give access to a large class of states, they also have the advantage of providing a fairly explicit description of these states, in contrast to an approach where the quantum state space is directly defined as the set of *all* possible quantum states (like in AQFT [17, 18], which we mentioned in section 1.). In the latter case, quantum states have to be specified through the expectations values they prescribe for arbitrary products of observables⁹. The difficulty is that there are relations that these expectations values together must satisfy (in particular to ensure that all measurement probabilities will come out positive, see [17, part III, def. 2.2.8]) making it non-trivial to actually construct valid states in this form and study their properties. Arguably, quantum states expressed as collections of partial density matrices are also subjected to

⁹ See however an alternative formulation, advocated in [21], allowing to recast the full state space of AQFT in projective form.

the satisfaction of extra relations (namely, the consistency condition from eq. (3)). Yet, with an additional assumption on the *size* of the label set, it is possible to obtain an explicit parametrization of all available projective states [30, subsection 2.1].

Suppose first that our partial theories are labeled by an increasing sequence: say, for concreteness, $(\eta_n)_{n \geq 1}$, where the partial theory labeled by η_n consists of lattice dof. on a regular lattice with step $a/2^n$. Then, the full quantum state space can be constructed very easily: indeed, each projective state can be constructed via a recursive process, starting from a density matrix ρ_{η_1} on \mathcal{H}_{η_1} , and choosing, for any $n \geq 1$, $\rho_{\eta_{n+1}}$ in the pre-image $\text{Tr}_{\eta_{n+1} \rightarrow \eta_n}^{-1} \langle \rho_{\eta_n} \rangle$ (taking advantage of the fact that the partial trace is always surjective).

The class of label sets for which a *constructive* description of the quantum state space can be obtained in this manner is actually much broader. This follows from the following observation. If we have a projective state space build on a label set \mathcal{L} , then restricting it to a subset \mathcal{K} does *not* change the quantum state space, provided \mathcal{K} is *cofinal* in \mathcal{L} , ie. that it satisfies the condition

$$\forall \eta \in \mathcal{L}, \exists \kappa \in \mathcal{K} / \eta \preceq \kappa.$$

The projective state spaces are the same because coarser partial states can always be reconstructed from finer ones: if we know ρ_κ for all $\kappa \in \mathcal{K}$, we can reconstruct any ρ_η as $\text{Tr}_{\kappa \rightarrow \eta} \rho_\kappa$ for some $\kappa \succeq \eta$ [28, prop. 2.6]. For example, if we had chosen the sequence of lattice steps above to be $a/4^k, k \geq 1$, we would in fact have been describing the *same* state space, since $\{2k \mid k \geq 1\}$ is a cofinal part of $\{n \mid n \geq 1\}$.

Now, any *countable* label set admits a cofinal sequence of increasingly fine labels (this follows easily from the directedness property mentioned in subsection 2.1, see [30, prop. 2.1]). Furthermore, the same observation also implies that it is not really the *cardinality* of the label set that matters, but rather its *cofinality* (the smallest possible cardinality of a cofinal part in it). To summarize, constructive descriptions are readily available for all projective quantum state spaces built on label sets having countable cofinality.

3.3 Projective State Spaces Are Robust With Respect to the Selection of Partial Theories

As each label extracts finitely many dof. out of the full theory, expecting the label set to be of countable cofinality effectively requires that the full theory be spanned by countably many basic observables. This is a physically legitimate assumption, since it should be possible to determine and specify *which* observable has been measured in a given experiment using only a finite amount of information. This does not forbid the classical field theory to be formulated in terms of a continuum of dof. for mathematical convenience, but we should keep in mind that infinitesimally close observables will be indistinguishable in practice, so that a dense and countable subset of basic observables should be sufficient to capture the entire physical content of the theory.

Nevertheless, one might be worried about making the construction of the quantum state space dependent on the selection of a more or less arbitrary subset of observables. In particular, it seems that the *universality* put forward in section 2. would suffer. Continuing the example started in subsection 3.2 above, we could have chosen the sequence of lattice steps to be, say, $a/3^p, p \geq 1$. In the absence of a background geometry, the choice of a sequence of finer and finer lattices would involve an even greater degree of arbitrariness, since there is no notion of 'regular lattices' in this case (see the construction in [34, section 3] and [26, section II]). The natural question to ask is therefore whether these various options lead to dramatically different quantum state spaces.

Remarkably, there is a simple condition that an increasing sequence of labels should satisfy to ensure that the projective state space built on it will be universal [30, subsection 2.2]. Denote by $\mathcal{L}^{(\text{ext})}$ the complete set of all possible labels, among which a sequence has to be chosen. As explained above, if a cofinal sequence $(\kappa_k)_{k \geq 1}$ can be found in $\mathcal{L}^{(\text{ext})}$, the projective state space built on it coincides with the one that could be built on $\mathcal{L}^{(\text{ext})}$ and is thus universal. But if $\mathcal{L}^{(\text{ext})}$ is simply too big to admit such a cofinal sequence, we can relax the condition and only require the sequence $(\kappa_k)_{k \geq 1}$ to be *quasi-cofinal* in the following sense: for any label η in $\mathcal{L}^{(\text{ext})}$, it should be possible to slightly deform η so as to make it a sublabel of one of the κ_k . More precisely, it should be possible [30, def. 2.7]:

1. to let the deformation be arbitrarily small;
2. and, if some of the dof. contained in η are among the dof. captured by the sequence $(\kappa_k)_{k \geq 1}$ (namely, if there exists a κ_{k_0} that happens to be a sublabel of η), to demand that these particular dof. be left untouched by the deformation [30, fig. 2.1].

In the case of 1-dimensional lattices (and by extension of square lattices in any dimension), it transpires from the analysis in [30, section 3] that a sequence of finer and finer lattices will be quasi-cofinal if, and only if, the set of lattice nodes becomes dense as $k \rightarrow \infty$. In the case of free quantum fields on a possibly curved background (using the quantization scheme mentioned at the end of subsection 2.2), quasi-cofinal sequences of labels were found in [25, subsection 3.2] to correspond to dense, countably-generated vector subspaces in the infinite-dimensional classical phase space. Both results are in agreement with the physical discussion at the beginning of the present subsection: they demonstrate how dense subsets of dof. end up being sufficient to capture the physical content of the theory.

Let us clarify what we precisely mean by saying that the projective state space on a quasi-cofinal sequence is 'universal' [30, theorem 2.8]. If we have two different quasi-cofinal sequences, the set of dof. captured by the first one will of course not be *exactly* the same as the set of dof. captured by the second one. However, an arbitrarily small deformation can map the first set into the second one, and, *modulo* this approximate identification of the observables, the two quantum state spaces will coincide (ie. any state in the first one will be in *one-to-one* correspondence with a state in the second one, in such a way that the expectations values of the suitably identified observables match). An important corollary of this property [30, prop. 2.9] is that symmetries of the original field theory can be represented as isomorphisms of the thus constructed quantum state space: indeed, a

symmetry transformation, mapping a sequence (κ_k) to a sequence (λ_l) , can be composed with a small deformation mapping (λ_l) back to (κ_k) , to obtain a transformation that stabilizes the state space and observable algebra on (κ_k) .

Note that this notion of universality should not be confused with the one expressed by Fell's theorem [9]. Fell's theorem is a result stating that the Hilbert space \mathcal{H}_Ω constructed from a vacuum state Ω as discussed in section 1. is approximately universal in the following sense: for any quantum state Ω' that can be built on the same algebra of observables as Ω , any finite set of observables f_1, \dots, f_N and any desired precision ϵ , there exists a density matrix ρ on \mathcal{H}_Ω such that the expectation values of these *particular* observables, as computed with respect to Ω' vs. ρ , will not differ by more than ϵ . The most crucial difference between this result and the one discussed in the present subsection is that, if one want to fully take advantage of Fell's theorem to recover from \mathcal{H}_Ω a sufficiently universal quantum state space, one will effectively have to work, not in the space of density matrices on \mathcal{H}_Ω , but in a suitable *completion* thereof (namely, the completion with respect to the topology implied by the just spelled out statement of the theorem). By contrast, a projective state space constructed on a quasi-cofinal sequence is already 'complete' enough: its sole limitation lies in the fact that the quantum states it describes *a priori* only prescribe the expectations values of the observables in a certain countably-generated subset, safe in the knowledge that all other observables included in the (artificial but technically convenient) continuum classical field theory can be reconstructed from those at an arbitrary level of precision.

4. From Kinematical to Dynamical State Spaces

4.1 The Need for Regularization

We have not, so far, addressed the question of how the equations of motion of the original field theory are to be implemented on a quantum state space constructed along the lines of section 2.. This question raises two distinct kinds of issues: first, we need to understand how the dynamics can be expressed in a *classical* theory which has been reformulated as a collection of partial theories as described in subsection 2.1; second, we need to suitably quantize this classical dynamics. The second part of the problem will involve the same tools as form usually part of a canonical quantization endeavor [7, 14, 13], however the first part is specific to projective state spaces. The core difficulty is that the dynamics is unlikely to respect the coarse graining introduced in subsection 2.1. In a field theory we cannot, in general, write *closed equations* involving only *finitely* many degrees of freedom: the evolution of an infinite-dimensional system, when projected on a finite-dimensional truncation, is likely to be non-deterministic, because it is subject to back-reaction from all the discarded dof. whose actual value is unknown.

For the rest of the present section, we will adopt a parametrized description of the dynamics: by adding the time variable explicitly to the configuration space of our theory, solutions of the equations of motion can be represented as orbits or trajectories in the




Fig. 3 Adapted (left side) versus regularized dynamics (right side)

resulting extended phase space (see [45, section 1.8] as well as the procedure illustrated in [24, 23]). This way of looking at dynamics has the advantage that it generalizes seamlessly to gauge theories (where the solutions correspond to orbits under both time evolution and gauge transformations), and even to theories like general relativity, that do not refer to any external time variable (where time evolution is just another gauge transformation [1]). In this language, the dynamics of the original field theory are said to be *adapted* to a collection of partial theories $(\mathcal{M}_\eta)_\eta$ if they induce consistent partial dynamics on each \mathcal{M}_η [27, subsection 3.1]. Consistency here means that any orbit of the partial dynamics on a finer label η' should project exactly to an orbit on a coarser label η . In this way, we can obtain projective families of orbits, which correspond precisely to the solutions of the full dynamics (left part of fig. 3). As stressed above, we cannot expect this ideal case to apply in any realistic field theory (unless we dispose *beforehand* of an extensive understanding of the dynamics, that we can take into account when defining the collection of partial theories; this would suffer from the same limitations as underlined in section 1. in the context of vacuum-based quantization).

What we can do instead [27, subsection 3.2], is to define, for each η , *approximate* partial dynamics, that *can* be formulated in closed form over \mathcal{M}_η at the cost of not *exactly* matching the desired full dynamics.¹⁰ These approximate dynamics should become closer and closer to the exact one as we consider finer and finer labels η . This implies that such a collection of approximate partial dynamics will *not* be consistent in the sense above, because the partial dynamics on a finer label η' are more accurate than the ones on a coarser label η . Still, given some initial data $x(t=0)$ (or, more generally, given a point on a gauge fixing surface in \mathcal{M}_∞ that intersects each orbit exactly once), we can consider the family of approximate orbits $(Y_\eta)_\eta$, where Y_η denotes the orbit on \mathcal{M}_η originating from (the projection of) x (right part of fig. 3).

¹⁰ The approach taken in [27, subsection 3.2] is a slight generalization of this proposal, in which additional parameters can be introduced to control the approximation. This is done by indexing the approximation scheme on a larger label set, namely a certain part of $\mathcal{L} \times \mathcal{E}$, with the set \mathcal{E} holding the extra approximation parameters.

The idea behind this approach is to represent an exact solution of the full dynamics by the family of approximate solutions that converges to it. Of course, not all families of successive approximations will converge to an exact solution: the initial data x will in general have to satisfy suitable regularity conditions for the approximation scheme to converge. But, importantly, the convergence of a given family can be studied completely within the projective setup, *without* referring to the continuous theory \mathcal{M}_∞ . Let us define, for any $\eta \preceq \varepsilon$

$$Y_\eta^\varepsilon := \pi_{\varepsilon \rightarrow \eta} \langle Y_\varepsilon \rangle$$

ie. the orbit at a finer level of approximation ε , projected back on \mathcal{M}_η . As stressed above the fact that the partial dynamics defined on \mathcal{M}_η is only a coarse approximation is reflected by the deviation $Y_\eta \neq Y_\eta^\varepsilon$. Convergence of the approximation is then expressed as the requirement that, for each η , the orbit Y_η^ε converges, as ε gets finer and finer, to (the projection of) an exact orbit Y_η^∞ . Now, the family $(Y_\eta^\infty)_\eta$ is a projective family of orbits by construction, and as in the case of 'adapted dynamics' discussed above, it can be seen as the projective description of an exact solution of the full dynamics.

4.2 A Framework for Studying Convergence at the Quantum Level

The ability to discuss the convergence of the regularization scheme from within the projective setup is especially crucial when going over to the quantum theory, where there is no obvious equivalent of \mathcal{M}_∞ . Specifically, we will implement the dynamics on the quantum projective state space by turning the approximate partial dynamics on each \mathcal{M}_η into a corresponding quantum dynamics on \mathcal{H}_η , constructing for any given initial data an associated family of approximate solutions, and study its convergence with the same strategy as on the classical side. Note that a difference between classical and quantum dynamics is that solutions of the (approximate) quantum dynamics are just special quantum states, not *subsets* of states as in the classical case: this effect can be observed already at the level of classical *statistical* mechanics, where 'kinematical' states would be arbitrary probability distributions on the (extended) phase space, while solutions, or 'dynamical' states, would be those probability distributions that are *constant* along the orbits. Accordingly, a converging family of approximated solutions of the quantum dynamics will define a projective family of quantum states $(\rho_\eta^\infty)_\eta$; in other words, such a family will provide a special state in the kinematical projective quantum state space.

In [29, subsection 3.2], this procedure has been applied to the second quantization of the Schrödinger equation (aka. non-relativistic quantum field theory [6]). One-particle quantum mechanics, seen as a classical field theory, can be quantized in a traditional way leading to a Fock space which is simply the Fock space describing an arbitrary number of non-interacting particles, each obeying the first-quantized theory. Alternatively, a projective quantum state space can be built (the partial theories being obtained by orthogonally projecting the wavefunction on every finite dimensional subspace of the one-particle Hilbert space) and the dynamics can be regularized on it. Interestingly, a domain of convergence for this regularization of the quantum dynamics can be delimited

which is consistent with the traditional approach, in the sense that the dynamical projective quantum states it describes are precisely the ones in the Fock sector. This sheds light as to why the Fock space is an appropriate arena for describing this *non-interacting* theory and demonstrates how projective techniques can allow us to get a handle on the quantum dynamics without having to guess *beforehand* a good vacuum: instead, a sector on which the dynamics can be well-defined should naturally emerge from the study of the convergence behavior of approximate solutions. This also raises the intriguing open question of whether suitable domains of convergence will always take the form of a vacuum sector with respect to a certain (dynamics-aware) vacuum.

5. Outlook: Bridging Canonical Quantization and Path-integral Approaches

Looking for a good quantum state space may seem like a preoccupation only relevant if we insist to rely on canonical quantization. The argument here would be that we do not actually need any quantum state space: it is sufficient to know the quantum probabilities for the outcome of any arbitrary experimental protocol, and these probability distributions could themselves be reconstructed from their moments. However, the so-called *time-ordered* expectation values that can be computed from a path-integral formulation are *not* the moments corresponding to the successive measurement of the corresponding observables, notwithstanding the fact that the measurement operations are, of course, causally ordered. For example, to know the probability of measuring successively $f_1 = a_1, \dots, f_n = a_n$, we need to compute¹¹

$$\begin{aligned} P(f_1 = a_1, \dots, f_n = a_n) &= \\ &= \frac{\|\Pi_1 \Omega\|}{\|\Omega\|} \frac{\|\Pi_2 \Pi_1 \Omega\|}{\|\Pi_1 \Omega\|} \cdots \frac{\|\Pi_n \dots \Pi_1 \Omega\|}{\|\Pi_{n-1} \dots \Pi_1 \Omega\|} = \\ &= \sqrt{\langle \Omega | \Pi_1 \dots \Pi_n \Pi_n \dots \Pi_1 | \Omega \rangle} \quad (5) \end{aligned}$$

where Π_i denotes the spectral projector on the eigenvalue a_i of \hat{f}_i . Knowing all time-ordered expectation values is thus *a priori* not sufficient to recover the physically interesting moments (note the mirroring in the formula above: from right to left, the observables appear first in causal-order and are then repeated in *reverse* causal order). Reciprocally, it turns out that if we had the vacuum-vacuum expectation values for product of operators in arbitrary orders, which would encode all the physically interesting information, we could in fact also reconstruct the quantum state space, via the GNS construction mentioned in section 1. (in this context also known as Wightman reconstruction [44]).

Now, on Minkowski spacetime, the expectation values for products of observables in arbitrary order can *all* be reconstructed from the time-ordered ones (or, equivalently, the quantum expectation values for the different orderings of operators can be all reconstructed via analytic expansion from the *Euclidean* path-integral: different orderings

¹¹ This is assuming the universe was initially in the vacuum state, but we can easily relax this questionable assumption by considering relative probabilities, eg. $P(f_1=a_1, \dots, f_n=a_n)/P(f_1=a_1)$.

then arise as different prescriptions for going around singularities, as was proved by Osterwalder and Schrader [35, 36]). It is however not clear how the fairly subtle complex analysis underlying this result could be generalized to non-static background spacetimes (although results exist when a timelike Killing vector is available [19]), or even to background-independent theories. This is the situation where disposing of a (kinematical) quantum state space can be valuable even for the path-integral approach. We can then, instead of only computing vacuum-vacuum expectation values, develop the path-integral between *arbitrary* boundary quantum states (like is done eg. in the spin foam approach [37] to quantum gravity, leaning against the Ashtekar-Lewandowski Hilbert space mentioned in section 1.) and this is indeed an alternative way of capturing all the physically interesting information (since we can then insert a decomposition of the identity in the middle of the above formula).

Projective quantum state spaces seem particularly well-suited to play the role of supportive state spaces for boundary-aware path-integral formalisms. On one hand, they implement the same kind of coarse-graining/refining which is used in practice to compute, through a limiting process, path-integral expressions. On the other hand, they anticipate the loss of privileged status for the vacuum state that is the natural consequence of allowing non-vacuum boundaries. Reciprocally, techniques that have been developed in the context of path-integral approaches could be applied to the strategy presented in section 4., to help with the study of *convergence*, and with the inclusion of *renormalization* (which was not needed for the simple example considered in [29, subsection 3.2] but is likely to come into play when considering interacting theories).

Acknowledgement

I would like to thank Thomas Thiemann for encouraging me to write this overview letter.

Financial support through the Programa de Becas Posdoctorales de la UNAM is acknowledged.

References

- [1] R. ARNOWITT, S. DESER, AND C. W. MISNER, *The Dynamics of General Relativity*, in Gravitation: an Introduction to Current Research, L. Witten, ed., Wiley, 1962, ch. 7, pp. 227–265, arXiv:gr-qc/0405109.
- [2] A. ASHTEKAR, *Lectures on Non-perturbative Canonical Gravity*, no. 6 in Advanced Series in Astrophysics and Cosmology, World Scientific, 1991. Notes prepared in collaboration with Ranjeet S. Tate.
- [3] A. ASHTEKAR AND C. ISHAM, *Representations of the Holonomy Algebras of Gravity and non-Abelian Gauge Theories*, Class. Quant. Grav., 9 (1992), p. 1433, arXiv:hep-th/9202053.
- [4] A. ASHTEKAR AND J. LEWANDOWSKI, *Representation Theory of Analytic Holonomy C^* -Algebras*, in Knots and Quantum Gravity, J. Baez, ed., Oxford University Press, 1994, pp. 21–62, arXiv:gr-qc/9311010.

-
- [5] J. F. BARBERO G., T. PAWŁOWSKI, AND E. J. S. VILLASEÑOR, *Separable Hilbert Space for Loop Quantization*, Phys. Rev. D, 90 (2014), p. 067505, arXiv:1403.2974.
- [6] P. A. M. DIRAC, *The Quantum Theory of the Emission and Absorption of Radiation*, Proc. Roy. Soc., A114 (1927), p. 243.
- [7] P. A. M. DIRAC, *Lectures on quantum mechanics*, no. 2 in Belfer Graduate School of Science Monographs, Academic Press, 1964.
- [8] W. FAIRBAIRN AND C. ROVELLI, *Separable Hilbert space in Loop Quantum Gravity*, J. Math. Phys., 45 (2004), pp. 2802–2814, arXiv:gr-qc/0403047.
- [9] J. M. G. FELL, *The Dual Spaces of C^* -Algebras*, Trans. Amer. Math. Soc., 94 (1960), pp. 365–403.
- [10] C. FLEISCHHACK, *Representations of the Weyl Algebra in Quantum Geometry*, Commun. Math. Phys., 285 (2009), pp. 67–140, arXiv:math-ph/0407006.
- [11] I. GELFAND AND M. NAIMARK, *On the Imbedding of Normed Rings into the Ring of Operators in Hilbert Space*, Mat. Sb., Nov. Ser., 12 (1943), pp. 197–213.
- [12] K. GIESEL AND T. THIEMANN, *Algebraic Quantum Gravity (AQG) I. Conceptual Setup*, Class. Quant. Grav., 24 (2007), pp. 2465–2498, arXiv:gr-qc/0607099.
- [13] D. GIULINI AND D. MAROLF, *A uniqueness Theorem for Constraint Quantization*, Class. Quant. Grav., 16 (1999), pp. 2489–2505, arXiv:gr-qc/9902045.
- [14] ———, *On the Generality of Refined Algebraic Quantization*, Class. Quant. Grav., 16 (1999), pp. 2479–2488, arXiv:gr-qc/9812024.
- [15] L. GÅRDING AND A. WIGHTMAN, *Representations of the Commutation Relations*, Proc. Natl. Acad. Sci. USA, 40 (1954), pp. 622–626.
- [16] R. HAAG, *On Quantum Field Theories*, Dan. Mat. Fys. Medd., 29 (1955), pp. 1–37.
- [17] ———, *Local Quantum Physics – Fields, Particles, Algebras*, Texts and Monographs in Physics, Springer, second ed., 1996.
- [18] S. HOLLANDS AND R. M. WALD, *Quantum fields in curved spacetime*, Physics Reports, 574 (2015), pp. 1 – 35, arXiv:1401.2026.
- [19] A. JAFFE AND G. RITTER, *Quantum field theory on curved backgrounds. I. The Euclidean functional integral*, Commun. Math. Phys., 270 (2007), pp. 545–572, arXiv:hep-th/0609003.
- [20] J. KIJOWSKI, *Symplectic Geometry and Second Quantization*, Rep. Math. Phys., 11 (1977), pp. 97–109.
- [21] J. KIJOWSKI AND A. OKOLOW, *A modification of the projective construction of quantum states for field theories*, J. Math. Phys., 58 (2017), p. 062303, arXiv:1605.06306.
- [22] T. KOSŁOWSKI AND H. SAHLMANN, *Loop Quantum Gravity Vacuum with Nondegenerate Geometry*, SIGMA, 8 (2012), arXiv:1109.4688.
- [23] K. KUCHAR, *Dirac Constraint Quantization of a Parametrized Field Theory by Anomaly-free Operator Representations of Spacetime Diffeomorphisms*, Phys. Rev. D, 39 (1989), pp. 2263–2280.
- [24] ———, *Parametrized Scalar Field on $\mathbb{R} \times S^1$: Dynamical Pictures, Spacetime Diffeomorphisms, and Conformal Isometries*, Phys. Rev. D, 39 (1989), pp. 1579–1593.

-
- [25] S. LANÉRY, *Polarization-free Quantization of Linear Field Theories*. Preprint, 2016, arXiv:1610.08740.
- [26] S. LANÉRY AND T. THIEMANN, *Projective Loop Quantum Gravity I. State Space*, J. Math. Phys., 57 (2016), p. 122304, arXiv:1411.3592.
- [27] —, *Projective Limits of State Spaces I. Classical Formalism*, J. Geometry Phys., 111 (2017), pp. 6–39, arXiv:1411.3589.
- [28] —, *Projective Limits of State Spaces II. Quantum Formalism*, J. Geometry Phys., 116 (2017), pp. 10–51, arXiv:1411.3590.
- [29] —, *Projective Limits of State Spaces III. Toy-Models*, J. Geometry Phys., 123 (2018), arXiv:1411.3591.
- [30] —, *Projective Limits of State Spaces IV. Fractal Label Sets*, J. Geometry Phys., 123 (2018), arXiv:1510.01926.
- [31] J. LEWANDOWSKI, A. OKOŁÓW, H. SAHLMANN, AND T. THIEMANN, *Uniqueness of Diffeomorphism Invariant States on Holonomy-Flux Algebras*, Commun. Math. Phys., 267 (2006), pp. 703–733, arXiv:gr-qc/0504147.
- [32] A. OKOŁÓW, *Quantization of Diffeomorphism Invariant Theories of Connections with a Non-compact Structure Group – an example*, Comm. Math. Phys., 289 (2009), pp. 335–382, arXiv:gr-qc/0605138.
- [33] —, *Construction of Spaces of Kinematic Quantum States for Field Theories via Projective Techniques*, Class. Quant. Grav., 30 (2013), p. 195003, arXiv:1304.6330.
- [34] —, *Kinematic Quantum States for the Teleparallel Equivalent of General Relativity*, Gen. Rel. Grav., 46 (2014), p. 1653, arXiv:1304.6492.
- [35] K. OSTERWALDER AND R. SCHRADER, *Axioms for Euclidean Green’s functions*, Comm. Math. Phys., 31 (1973), pp. 83–112.
- [36] —, *Axioms for Euclidean Green’s functions. II*, Comm. Math. Phys., 42 (1975), pp. 281–305.
- [37] C. ROVELLI, *Zakopane Lectures on Loop Gravity*. Preprint, 2011, arXiv:1102.3660.
- [38] I. SEGAL, *Irreducible Representations of Operator Algebras*, Bull. Amer. Math. Soc., 53 (1947), pp. 73–88.
- [39] M. H. STONE, *On One-Parameter Unitary Groups in Hilbert Space*, Ann. Math., 33 (1932), pp. 643–648.
- [40] T. THIEMANN, *Modern Canonical Quantum General Relativity*, Cambridge Monographs on Mathematical Physics, Cambridge University Press, 2007.
- [41] J. VON NEUMANN, *Die Eindeutigkeit der Schrödingerschen Operatoren*, Math. Ann., 104 (1931), pp. 570–578.
- [42] —, *Über Einen Satz Von Herrn M. H. Stone*, Ann. Math., 33 (1932), pp. 567–573.
- [43] —, *On Infinite Direct Products*, Compos. Math., 6 (1939), pp. 1–77.
- [44] A. S. WIGHTMAN, *Quantum Field Theory in Terms of Vacuum Expectation Values*, Phys. Rev., 101 (1956), pp. 860–866.
- [45] N. M. J. WOODHOUSE, *Geometric Quantization*, Oxford Mathematical Monographs, Oxford University Press, second ed., 1992.

