

Referee Report

Causality and realizability of local operations in quantum field theory

arXiv:2607.12976

Summary and overall assessment

The manuscript studies the relation between two proposed restrictions on local operations in algebraic quantum field theory. The first is the past-support-non-increasing (PSNI) condition used to exclude Sorkin-type superluminal signaling. The second is realizability in the Fewster–Verch (FV) probe-coupling framework. The authors introduce finite Weyl-algebraic POVMs and instruments, single out the subclass whose nonselective channel is a convex combination of Weyl displacements, and call these instruments diagonal. They then claim: (i) strong-density results for Weyl POVMs and normal instruments; (ii) a characterization of causal Weyl instruments and closure under communication-network composition; (iii) genericity of the diagonal subclass; (iv) asymptotic FV realization of every diagonal Weyl instrument and hence nondemolition implementation of arbitrary local POVMs; (v) causal “box channels” encoding arbitrary nonsignaling correlations, including examples that cannot be approximated by FV schemes; and (vi) undecidability of a gapped network-implementability problem by reduction from the $\text{MIP}^*=\text{RE}$ characterization of quantum correlations.

This is an ambitious and potentially valuable program. In particular, the proposed bridge between relativistic measurement theory and the distinction between nonsignaling, commuting-operator, and tensor-product correlations is conceptually attractive. If correct, a concrete separation between PSNI causality and FV realizability would be important for the foundations of QFT measurement theory. The diagonal-instrument construction could also provide a useful positive counterpart to existing impossible-measurement results. The paper addresses current literature directly: the FV framework, the causal analyses of Jubb and Oeckl, recent local-POVM preparation protocols, the finite-dimensional distinction between causal and localizable channels, and $\text{MIP}^*=\text{RE}$ all enter in a natural way.

Unfortunately, the main conclusions are not established by the present manuscript. There are substantive gaps in several indispensable links of the argument, as well as many equations that are internally inconsistent. The problems are not confined to exposition. They affect the locality lemma used in the causal classification, the construction asserted to be an FV scheme, the recovery of a nonsignaling box from a QFT channel, the examples intended to prove the unbounded-interaction-zone property, and the network theorem used in the undecidability reduction. Because the headline conclusions depend on these links in sequence, correcting isolated typographical errors would not suffice.

Recommendation

I recommend rejection in the present form. A substantially reconstructed manuscript could merit fresh consideration, because the core questions and several of the proposed constructions are interesting. Before that stage, however, the authors need to replace the informal or inconsistent arguments by complete proofs, state the representation-theoretic hypotheses precisely, and verify

every displayed formula in the central constructions. In its present state the paper does not meet the correctness standard required for publication in a journal such as PRD or JHEP.

Major comments

1. The representation and locality assumptions are overstated, and Lemma 13 is not proved.

Section II B says that the GNS vector of an arbitrary algebraic state has the Reeh–Schlieder property. Cyclicity for the global algebra is automatic in the GNS construction; cyclicity for every local algebra is not. The paper later uses local cyclicity (and effectively local faithfulness) in the probe construction, for example to conclude that the projection in equation (77) has nonzero expectation. The admissible class of states and representations must therefore be specified at the outset. Likewise, the Haag property asserted in Section II C is a strong hypothesis whose validity can depend on topology, the field equation, and the choice of net; it should not be presented as an automatic consequence of the three elementary AQFT axioms.

More importantly, the proof of Lemma 13 in Appendix D does not justify its key inference. Proposition 46 shows, under suitable Haag assumptions, that actual Kraus operators of a local normal instrument belong to an appropriate local algebra. The manuscript then chooses a Weyl matrix representation with nonzero diagonal entries and concludes that each individual exponential e^{iQ_j} commutes with the causal-complement algebra. This does not follow. A positive matrix can have every diagonal entry nonzero while its range is a proper subspace; a rank-one matrix proportional to $(1,1)(1,1)^*$ is the simplest example. Locality of the Kraus combination $e^{iQ_1} + e^{iQ_2}$ does not by itself localize its two summands. The subsequent statement that commutation implies $\text{supp } f_j \subset O$ also confuses support of a representative with localization of its equivalence class or support of Ef_j .

Lemma 13 feeds directly into Theorem 17, Theorem 20, and parts of Propositions 22 and 34. The authors need either a new proof that extracts a matrix representation whose Weyl terms are individually localizable, or a reformulation of all later results in terms of the local Kraus span. Proposition 22 also invokes this reasoning for arbitrary normal channels localized in U^A and U^B , although Lemma 13 is stated only for Weyl instruments. The precise additional bimodule/Haag hypothesis needed for local Kraus representatives should be stated and proved.

There are related elementary inconsistencies that should have been caught before submission. In equation (26),

$$e^{i\Phi(f)}e^{i\Phi(g)}e^{-i\Phi(f)} = e^{-iE(f,g)}e^{i\Phi(g)},$$

not $e^{-iE(f,g)/2}e^{i\Phi(g)}$. Equation (41) ends with 0.1 although normalization requires zero. These errors are particularly concerning because the same Weyl phases and normalization identities are used throughout the paper.

2. The causal characterization and network-composition proofs require a rigorous reconstruction.

The statement of Theorem 17 quantifies over a “nonzero $d \in \mathcal{S}(d)$,” which is not meaningful; the intended set appears to be $\mathcal{D}$. Appendix E contains further mismatches between the statement and proof: d is alternately a test function and an equivalence class, J^- is replaced by J in one branch, and equation (E.12) drops the Weyl factor $e^{i\Phi(h+d)}$. Some of these may be transcription errors, but they occur inside the central necessary-and-sufficient criterion.

The geometric lemmas also need their missing hypotheses. Proposition 45 begins by shrinking R_- to a relatively compact region R'_- while preserving the domain-of-dependence inclusion, without proving that this is possible. The proof uses an undefined operation $H(\tilde{T}_+)$. If boundedness or compact containment of all PSNI test regions is required, it must be stated. Lemma 47 similarly moves between a representative supported in O , a class whose causal solution is supported in $J(O)$, and localizability in R_- without establishing the equivalence used in the theorem.

The proof of Theorem 20 in Appendix F is not presently checkable. The recursion contains undefined or inconsistent index sets, the induction statement does not match the cases used to update it, and one of the decisive sentences in item 2 is incomplete. The argument also relies on the unproved form of Lemma 13 to localize selective Kraus operators. Since selective feed-forward is exactly where PSNI of the nonselective channel is insufficient, this step cannot be left implicit. I recommend stating a clean two-node theorem first, with all spacetime regions and outcome dependencies explicit, and only then proving the finite directed-acyclic-network result by induction.

3. The diagonal FV construction is not, as written, a sequence of valid FV schemes.

The positive result in Section V A is among the most interesting claims, but several foundational points are missing. Most directly, equation (86) explicitly says that ω'_λ is not normalized. A nonnormalized positive functional is not a probe state and does not define the unital conditional expectation required in equations (21)–(23). Saying that its norm tends to one does not make any finite- λ member a physical FV scheme. The authors should normalize it for every λ and carry the normalization factor through Proposition 21.

Equation (88) divides by $\sqrt{p_j p_k}$ although the preceding definition permits $p_j = 0$. Zero-weight rows can presumably be removed using positivity, but this must be done before the division. Equation (91) defines the failure effect by subtracting the first s effects, whereas the instrument has A original outcomes; the sum must run to A . Equation (92) has the target field mixing back into Φ in both terms. The calculation in Appendix H instead uses the required probe term $\lambda\Psi(g'_\lambda)$. Equations (92) and (H.20) are also mutually inconsistent with the exact rotation (93): direct substitution into (93) gives $\Psi(h_j)$ in (H.20), not $\cos(2\lambda)\Psi(h_j)$.

Beyond these formula errors, the modal factorization needs justification in the representation being used. The finite collection of field quadratures can be put into symplectic normal form, but the claims that equation (77) is a genuine projection of the displayed tensor-product form, that postselection produces the stated quasifree covariance for an arbitrary chosen representation, and that the displaced projections become orthogonal in the topology needed to infer $\mu^\lambda \rightarrow 1$ should be proved. The interaction quoted from the local swap construction must also be stated with the exact hypotheses under which (93) holds. Finally, the conclusion should consistently say “asymptotically FV-realizable.” It does not follow that photon counting or a general local POVM has an exact finite-energy nondemolition realization.

The status of the exact rotation in equation (93) is especially important: it is not a routine consequence of the FV axioms, and Appendix H assumes it rather than deriving it. The manuscript should either prove a precise compactly supported scattering theorem for the coupled target–probe equations or cite a source in which that theorem, with the required geometric and well-posedness hypotheses, is actually established. Without this input, even a corrected normalization calculation would not complete Proposition 21.

4. The box-channel recovery argument is only heuristic, so the separation from FV realizability is not yet established.

The box-channel idea in Section V B is appealing: measure approximate modular values, condition later displacements on a nonsignaling table, and recover that table in a Bell experiment. However, Proposition 26 is derived by repeatedly applying equations (118) and (124), both written only with $\ll$ and $\approx$. No error norm, rate, or uniformity under subsequent channels is supplied. Equation (140) is an exact four-delta limit, and the later non-FV and distance claims require precisely this exact result. The authors should replace the heuristic concentration discussion by a spectral-calculus estimate for the trigonometric Kraus operators, propagate the bounds through the finite composition, and prove convergence in the topology used by the definition of asymptotic realizability.

There are also algebraic mistakes in the discrete measurement itself. In the expansion following equation (115), the inner phase is printed as $e^{i2\pi\gamma(k-j)}$; without division by n it sums to n , rather than producing the Kronecker delta used on the next line. The existence of the squeezed states in Proposition 26 should be shown within the fixed folium relevant to the asymptotic FV definition, not merely asserted for a varying family of abstract quasifree states.

Proposition 22 is stated for arbitrary normal channels local to U^A and U^B , but its proof assumes local Kraus operators and cites a Weyl-specific lemma. The counterexample only needs the displacement channels later constructed, so the proposition could be restricted accordingly; otherwise a general local-Kraus theorem is required. The closedness argument involving C_{qc} should also spell out the two limits: first the FV approximation for each fixed squeezed state, then the squeezing limit.

5. The manuscript does not prove that the required box/Bell setups exist for the claimed scalar theories.

Proposition 28 is supported by Appendix I, but equation (I.4), the displayed causal propagator for the massless field in $1 + 1$ dimensions, subtracts an integrand from itself over the same domain and is therefore identically zero. All subsequent formulas (I.12)–(I.13) assume a nonzero propagator. The proof also continues to use uncontrolled approximate equalities when exact nonvanishing and linear-independence statements are required. A corrected causal kernel and explicit estimates in the smoothing parameter are necessary.

The generic-theory route in Proposition 29 has incompatible support assignments. Lemma 61 in Appendix J first declares the functions $\hat{g}_{\hat{X}}^i$ to be supported in T^X , then constructs them with a bump supported in U^X . It later chooses f_X in U^X , although equation (134) and Proposition 26 require f_X in R_+^X . The restrictions in (J.4) are correspondingly taken on the wrong region. These are not harmless labels: the spacelike commutators and temporal order of the Bell experiment depend on them. The final list of pairwise causal-separation relations also does not visibly cover every mixed A/B pairing in equation (134).

The genericity results in Appendix G need additional functional analysis. The density proof switches to an $L^2(V)$ norm even when V is not relatively compact and never relates this norm to the test-function topology. The Baire argument uses sets that are shown dense but not shown open. The finite-rank wavefront-set argument in Lemma 58 and the passage from a null relation with a regular compact to one with its interior also need proof. These issues matter because genericity is used both to identify diagonal causal channels and to supply the examples underlying the main separation theorem.

6. The network-FV theorem and the undecidability reduction are not established.

Proposition 34 is the key bridge from implementable QFT channels to C_{qa} correlations. Appendix K currently contains mutually inconsistent definitions ($\vee_j$ versus $\wedge_j$, an index set initialized as $\{\emptyset\}$, and undefined quantities such as $c(P_j)$), and the factorization induction is not a proof in its present form. It calls Heisenberg-picture completely positive maps “trace preserving” rather than unital, assumes selective probe measurements can always be assigned to exactly one Bell wing, and again assumes local Kraus representations for arbitrary normal channels. A replacement proof should define the network as a finite causal directed acyclic graph, state the allowed initial correlations of all probes and ancillas, and derive a tensor-product strategy using the split inclusion. If the initial state can be nonnormal, the conclusion may be membership in the closure C_{qa} rather than in C_{qs} at the finite stage.

There is a concrete invalid step in this factorization argument: Appendix K replaces a selective Lueders expression $\sqrt{M_c} O_c \sqrt{M_c}$ by $M_c O_c$ without establishing that M_c commutes with the future observable O_c . The latter expression is not in general a completely positive operation. Since the processing region can be timelike related to the later readout, locality does not supply the missing commutation relation. This step must be replaced, not merely notationally corrected.

Proposition 36 also needs revision. Equation (154) uses Alice’s input in Bob’s POVM and has inconsistent outcome indices. More substantially, the definition of an FV-friendly box channel does not state the cross-separation conditions needed for the two ancillary modal algebras to be represented as independent tensor factors. The assertion that two local modes automatically generate $B(\hat{H})^{\otimes 2}$ must be tied to a regular representation and an explicit symplectic direct sum.

The convergence definitions are also inconsistent at this point. Definition 32 tests implementability against all states on $B(H)$, while Section II F tests asymptotic FV realizability in a GNS folium and Proposition 21 supplies strong-operator convergence. Uniformly bounded strong convergence controls expectations of normal states, not arbitrary singular states. Thus Proposition 36 does not establish implementability under the printed Definition 32. The authors should consistently restrict the definition to normal/folium states or prove a genuinely stronger convergence statement.

The final reduction has two further logical issues. Problem II is written as a promise problem: output YES for implementable channels and NO for channels that are not $1/(3M^2)$ -implementable. The finite net queried in the proof can contain channels in the gray zone between these cases. A promise oracle cannot be queried there. The theorem can perhaps be formulated as a total weak-membership problem whose algorithm must halt on every rational input and may answer arbitrarily in the gray zone; that definition, rather than the present table, is what the subsequent reasoning uses. The authors should also give a finite encoding of the fixed QFT geometry, test functions, and rational box so that “Turing machine input” is well defined.

Finally, Lemma 39 is true in spirit but not proved by Appendix M as printed. Equation (M.3) uses the facet half-space $B_j(P) \leq K_j$ inside C_l ; that intersection has the dimension of C_l , not one less. The facet is defined by equality. Equation (M.6) imposes $\sum_{a,b} \Delta P(a, b|0, 0) = 1$, contradicting the normalization-tangent condition (M.4), which requires zero. Thus the stated linear program is not the one whose rational optimum is needed in Theorem 38. These errors must be corrected before the computability claim can be assessed.

7. The Weyl positivity and approximation results need a precise C^* -algebraic setting.

Lemma 1 asserts that positivity of an algebraic Weyl element in one representation implies an SOS decomposition, arguing that simplicity makes the element positive in every representation. Simplicity by itself concerns kernels, not order, unless the authors are working in a specified simple C^* -completion and the representation extends faithfully to it. The manuscript instead defines $\widehat{\mathcal{W}}$ as an algebraic star algebra. The meaning of strict positivity $\pi(w) > 0$ is also important: a positive operator with trivial kernel need not be bounded below, whereas the Archimedean Positivstellensatz normally requires a uniform strict bound. The authors should specify the universal CCR C^* -algebra, regularity and nondegeneracy assumptions, and the exact theorem being applied.

Propositions 5 and 7 and Theorem 9 depend on this point. Their strong-operator approximation arguments are plausible but need a clean separation between nets and sequences, strong and strong-star convergence, and algebraic versus von Neumann closures. In Appendix C, the last normalization argument states $\hat{Z}_B = \check{Z}_B - \check{\check{Z}}_B$, whereas equations (C.18)–(C.19) give $\check{\check{Z}}_B = \hat{Z}_B + \check{Z}_B$. The desired conclusion may be recoverable from the diagonal trace calculation, but the sign and the proportional-channel normalization must be corrected explicitly.

More fundamentally, the sign/phase conjugations used after equation (C.8) preserve positivity of coefficient matrices but do not automatically preserve the cancellation identities (C.4)–(C.5) that make the associated operation unital up to scale. The authors must prove that every boundary term produced in Lemma 10 is proportional to a channel. Otherwise the decomposition required by Theorem 9, and hence an important input to the Section V A realization claim, remains unproved even after correcting the final sign.

Novelty and significance

The work could be significant if the central theorems survive revision. The manuscript does more than rephrase known Sorkin or FV results: it proposes an algebraically tractable class of operations, connects causal operations to nonsignaling boxes, and attempts to distinguish one-shot FV realizability from network implementation. The use of C_{qc} for a single FV coupling and C_{qa} for split network implementations is potentially insightful. The undecidability observation would also be novel in this setting if supported by a precise computational model.

The revised paper should sharpen the comparison with the cited literature. It should state exactly what is inherited from the FV scattering construction and recent local-POVM preparation protocol, what is new relative to the causal/localizable channel separation in quantum information, and which topology makes the nondemolition result stronger than existing asymptotic smeared-field measurements. Claims such as “any POVM can be realized” and “photon counting without perturbing the rest” should consistently include their asymptotic, representation-dependent, and potentially unbounded-resource qualifications.

Minor and presentation comments

The density of formula and notation errors is itself a barrier to evaluation. In addition to the issues above, I noted the following examples.

- The definition preceding Theorem 17 writes $Q_j = e^{-i\Phi(f_j)}$ even though Q_j is a quadrature, not a Weyl unitary.
- Equation (145) gives the CHSH witness threshold as $(2 + \sqrt{2})/2$. With the displayed coefficient $\eta = \delta_{a \oplus b, \alpha \beta}/4$, the commuting-operator bound is $(2 + \sqrt{2})/4$.
- Equation (146) contains $B(\Gamma, \Omega_P)$ although only $B(\Omega)$ is defined.
- Equation (159) is missing a closing parenthesis in the definition of μ .
- Appendix L ends by taking $\lambda \rightarrow \infty$, whereas Proposition 26 and the preceding equations use $\lambda \rightarrow 0$. The same appendix switches the box displacements from the f_X of equation (127) to g_X .
- Several displayed equations omit required sums or factors. For example, the first line of the proof of Proposition 37 suppresses the sums over a, b , and equation (E.12) suppresses both a range and an operator factor.
- The manuscript alternates among “local,” “localizable,” “interaction zone,” “coupling zone,” and “processing zone.” A compact glossary and consistent notation would greatly improve the argument.
- Numerous malformed words and encoding artifacts remain in the text, and several theorem statements contain mismatched variables or brackets. A complete line-by-line proofread is essential after the mathematical revision.

Conclusion

The paper asks the right questions and contains a promising conceptual synthesis, but its main results presently rest on unproved locality assertions, nonphysical intermediate “states,” heuristic limiting arguments, and internally inconsistent appendices. I would be interested in a future version that isolates a smaller set of theorems, proves each under explicit AQFT and representation hypotheses, and only then rebuilds the box-channel and undecidability consequences. The current version is not suitable for publication.