

Referee Report

Off-shell equivalence in quantum field theory and gravity

arXiv:2607.12644

Summary

The manuscript proposes an operational framework for deciding when two formulations of a classical or quantum field theory are equivalent beyond the usual on-shell equivalence theorem. The authors argue that equivalence must always be stated relative to a class of probes. They distinguish scalar kinematical readouts, classical response probes, scattering probes, and off-shell quantum probes. For tensorial observables, the proposed comparison uses contractions of configuration-space tensors with appropriately transported probe tensors. At the quantum level, the central object is the renormalized Vilkovisky–DeWitt (VDW) effective action, supplemented by a field-space metric and a matching of renormalized parameters.

The paper also distinguishes direct field redefinitions from *parent-mediated* equivalences. In the direct case, two descriptions live on locally diffeomorphic field spaces, and the action or VDW effective action is required to be related by pullback. In the parent case, both endpoints are obtained from a larger field space by reduction. The authors further divide these relations into local, branchwise, maximal-domain, and global equivalences. A scalar model with $\phi = \varphi + a\varphi^3/\Lambda^2$ illustrates the direct case and the failure of global invertibility. Metric $f(R)$ gravity, its multiplier formulation, its two-field scalar–tensor form, and the Jordan–Einstein map serve as the main gravitational example.

The two warm-up discussions are effective. Section 3 clearly separates the invariance of LSZ amplitudes from the change of a coordinate-linear susceptibility under a nonlinear field redefinition. Sections 6–8 also assemble, in one place, several distinctions that are often left implicit: the observable class, the field-space domain, the measure and source prescription, and the difference between a change of variables and an extension followed by reduction. These are worthwhile organizational ideas for the gravity and EFT communities.

The principal new content appears to be the attempt to combine this probe-relative taxonomy with parent reductions, branch gluing, and renormalized VDW representations. Unfortunately, the general parent-quantization argument is not correct as presently written, and the definition of parent-mediated equivalence is too weak to establish an equivalence between endpoints. In addition, the reduced measure used in the central argument is not defined invariantly. These are not peripheral presentation issues: they affect the main quantum claim and the $f(R)$ interpretation advertised in the abstract.

Recommendation

Major revision. The subject is appropriate for a journal such as JHEP or Physical Review D, and the proposed taxonomy could become useful. However, the central formal statements need substantial correction before the paper is suitable for publication. In particular, the authors must separate a delta-constrained lift of an endpoint theory from the ordinary quantization of an auxiliary-field parent, formulate the constrained measure coordinate-invariantly, and strengthen the

parent relation so that it actually identifies corresponding endpoint configurations and observables. The paper would also benefit greatly from one nontrivial quantum calculation, rather than an equality obtained after choosing the parent source and measure to reproduce the endpoint integral by construction.

Major comments

1. The delta-constrained construction is not the general quantization of an auxiliary-field parent.

The key step in Section 7.3 is the insertion

$$\delta[E_\Xi] \Delta_\Xi$$

in equations (7.24)–(7.28), where $E_\Xi = \delta\tilde{S}/\delta\Xi$. This localizes the path integral on a classical stationary section $\Xi = \Xi_\star(\Phi)$. It is a legitimate way to construct an endpoint-adapted constrained integral, but it is not what one normally obtains by quantizing the enlarged action $\tilde{S}[\Phi, \Xi]$ and integrating over an algebraic auxiliary field. The latter includes fluctuations in the Ξ direction and generally produces determinants and higher-order terms.

A finite-dimensional example makes the distinction immediate. Let

$$\tilde{S}(q, \xi) = S_0(q) + \frac{1}{2}A(q)\xi^2.$$

The ordinary integral over ξ produces a factor proportional to $A(q)^{-1/2}$, with the appropriate Lorentzian phase. By contrast,

$$\int d\xi \delta\left[\frac{\partial\tilde{S}}{\partial\xi}\right] \left|\frac{\partial^2\tilde{S}}{\partial\xi^2}\right| e^{i\tilde{S}/\hbar}$$

simply evaluates the integrand at $\xi = 0$ and has no such fluctuation factor. If $A(q)$ depends on the endpoint field, the difference cannot be discarded as an overall normalization. Thus equations (7.24)–(7.28) define a constrained theory chosen to reproduce a classical reduction; they do not prove that quantizing and eliminating an auxiliary field commute.

There is also a direct internal contradiction after equation (7.41). The manuscript states that $\tilde{S} + \Lambda^\alpha E_{\Xi, \alpha}$ leads to the same quantum theory as $\tilde{S}$. Integrating over Λ instead produces the delta-constrained theory, including the determinant or its ghost representation. This is precisely the construction that the surrounding paragraphs distinguish from the freely quantized parent in equations (7.42)–(7.43).

The authors should distinguish explicitly among: (i) direct endpoint quantization; (ii) an exact change of variables or Hubbard–Stratonovich identity; (iii) the endpoint-adapted delta-constrained lift introduced here; and (iv) unconstrained quantization of the enlarged action. The claim of a general parent equivalence should either be restricted to genuine multiplier parents, for which a delta functional follows from the action, or replaced by a theorem with hypotheses under which an exact functional identity is actually derived. The $f(R)$ multiplier identity in equation (8.45) is a special case of type (iii), not evidence for the unrestricted claim.

2. The reduced parent measure in equations (7.29)–(7.36) is not coordinate invariant.

The notation $s^*\tilde{\mu}$ is used for evaluating the full parent density on a lower-dimensional section in a chosen split (Φ, Ξ) . A top-dimensional density does not possess such a pullback to a lower-dimensional submanifold. Merely substituting $\Xi = \Xi_*(\Phi)$ into $\sqrt{\det \tilde{G}}$ depends on the chosen normal coordinates and on how the endpoint/auxiliary split is extended away from the section. This conflicts with the geometric covariance that motivates the use of the VDW formalism.

The statement following equation (7.36) that flat product metrics automatically satisfy measure compatibility is also false for a nonlinear section. Consider a flat two-dimensional parent with coordinates (q, ξ) , metric $dq^2 + d\xi^2$, and section $\xi = h(q)$. In adapted coordinates $\eta = \xi - h(q)$, the parent metric on $\eta = 0$ is

$$\begin{pmatrix} 1 + h'(q)^2 & h'(q) \\ h'(q) & 1 \end{pmatrix}, \quad \det \tilde{G} = 1.$$

The induced endpoint metric is $G_{\text{end}} = 1 + h'(q)^2$, so its Riemannian density is $\sqrt{1 + h'(q)^2}$, whereas evaluation of the ambient coordinate density gives 1. Equivalently, the Schur complement in equation (7.36) is $1/(1 + h'(q)^2)$, which is not constant in general. Thus flatness of the parent metric does not establish equations (7.32) or (7.36).

This section should be reformulated using an invariant coarea or constrained-measure formula. The geometrical delta, the norm or determinant of the constraint gradients, the induced metric on the constraint surface, and any orientation/absolute-value convention must be specified together. The authors should also show that the result is invariant under reparametrizations of the auxiliary coordinates and under changes of the local splitting that mix tangent and normal directions. Until this is done, the equality of generators in equation (7.37) is a coordinate-dependent assumption rather than a geometrical result.

The source term requires the same care. Pullback of the parent metric does not by itself imply that the parent geodesic displacement restricted to the section equals the geodesic displacement of the induced endpoint metric; the section would generally need an additional compatibility property, such as being totally geodesic, or the source must be defined intrinsically from the endpoint metric from the outset.

3. The parent-mediated classical criterion does not yet define an equivalence between the endpoints.

Equations (6.28)–(6.29) require each endpoint observable to be the pullback of some parent observable along its own section. They do not specify which point of endpoint 1 corresponds to which point of endpoint 2, nor do they require the two section images to coincide, intersect, or represent the same parent gauge orbit. Without such a correspondence, the condition is too weak: unrelated endpoint theories can be embedded into disjoint parts of a larger space, and their observables can be extended separately to parent functions.

The sentence after equation (6.29), which says that endpoint configurations are compared when their parent representatives describe the same configuration, is the needed condition, but it is not included in the definition. The authors should introduce an explicit relation, common reduction space, fiber product, or gauge-orbit identification that tells the reader when $s_1(\varphi_1)$ and $s_2(\varphi_2)$ correspond. Equality of the pulled-back actions in equation (6.30) is likewise not sufficient unless this cross-endpoint dictionary is specified.

This gap propagates to Section 7.6. The manuscript asserts that the union of the direct and parent criteria defines an equivalence relation and then forms the quotient in equations (7.76)–(7.77).

Reflexivity, symmetry, and especially transitivity are not established for different parents, local domains, branch choices, and accuracy-dependent parameter maps. The authors should prove closure under composition or common refinement. Otherwise the theory-space quotient should be presented as a proposed organizational picture rather than as a defined quotient set.

4. Covariant VDW tensors must be distinguished from physical response observables, and the in-in scope must be corrected.

Equations (5.5), (5.6), and (5.45) treat contractions of covariant derivatives of Γ^{VDW} with probe tensors as observables. Such contractions are scalars under field reparametrizations, but covariance alone does not make them measurable. A physical response requires a specified operator coupling, quantum state, contour, boundary conditions, and usually connected propagators obtained by inverting 1PI kernels. The Unruh–DeWitt example in equations (5.48)–(5.50), for example, uses a Wightman function; it is not derived from the in-out/Euclidean VDW effective action constructed in Section 4.

This matters because the manuscript repeatedly advertises retarded, Wightman, non-equilibrium, and in-in observables, while explicitly declining to construct a closed-time-path geometric effective action. The effective representation in equation (7.1) contains no state, contour, initial density matrix, or boundary-condition data. The authors should either develop the corresponding closed-time-path construction or restrict the claims to the Euclidean/in-out sector. A useful revision would work through one genuine detector or response observable in both parametrizations, including the transport of its state and source coupling.

There is also a tensorial point in equation (5.6). For rank three and above, successive covariant derivatives of a scalar are not generally fully symmetric once commutators act on its gradient. The geodesic Taylor series in equation (5.8) sees the fully symmetrized derivatives. If these tensors are intended to be covariant 1PI vertices paired with arbitrary rank- n probes, the authors should explicitly symmetrize the covariant derivatives or restrict the probes to the corresponding symmetric part.

Finally, the horizontal condition in equation (5.21) is not, by itself, sufficient to make the contraction with an arbitrary gauge-dependent theory tensor gauge invariant. The field dependence and gauge transformation of the horizontal projector or probe representative must be included, and the authors should state the conditions under which the tensor variation is purely vertical in each slot. A complete gauge-theory construction must also include the quotient or horizontal connection, orbit-volume factor, gauge fixing, and ghosts; the measure $\sqrt{\det G}$ on the redundant field space alone is not the gravitational VDW measure.

5. The status of the field-space metric is central and remains underdetermined in the main application.

Section 4.1 correctly acknowledges that the field-space metric is not uniquely fixed by the kinetic term in higher-derivative theories or EFTs. Nevertheless, the metric determines the connection, invariant measure, geodesic source, horizontal gauge projection, and the proposed off-shell response tensors. It is therefore not a harmless piece of notation. The claim in Section 4.2 that Bob infers the correct pullback metric automatically is true for the specially chosen one-component scalar example, but it does not provide a prescription for a generic theory.

The paper should state what fixes G for metric $f(R)$ gravity, for the (g, χ, λ) parent, and for the Jordan and Einstein descriptions. In Section 8.2 the equalities in equation (8.48) are imposed,

rather than derived from a definite gravitational field-space metric. Moreover, after integrating out λ , the section is written as $s_f(g) = (g, R[g])$, although the original parent had three variables. The status of the multiplier direction, its metric, and its contribution to the measure should be explicit.

If different admissible metrics define genuinely different off-shell quantum theories, the criterion is conditional on selecting that additional quantum input. If the authors instead intend the metric choices to be representational, they need a principle that identifies them. The abstract statement that one may choose an endpoint-adapted extension risks making the main equality tautological: the source and measure are selected precisely so that the parent integral reduces to the endpoint integral.

6. The branch analysis is incomplete for field configurations and the claimed maximal domains need a more precise formulation.

Equations (7.49)–(7.55) are the familiar finite-dimensional change-of-variables formula written functionally. For a pointwise nonlinear field map, however, an inverse-branch label can vary with spacetime position. In the cubic example with $a < 0$, the three uniform domains U_- , U_0 , and U_+ in equations (8.26)–(8.28) do not exhaust $\mathcal{F}_\varphi \setminus \mathcal{C}$. There are configurations whose values lie in different monotonicity intervals in different spacetime regions. A complete branch decomposition must say whether such configurations are admitted, what regularity and boundary conditions they obey, and how transitions between pointwise branches are treated. A formal sum over three global labels is not the full functional branch sum.

The paper also repeatedly speaks of the largest or the maximal equivalence domain. Non-injective maps can possess several incomparable maximal domains, with no unique largest one. The definitions should use maximal-by-inclusion domains, specify connected components or a branch selection, and avoid implying uniqueness without proof.

For the scalar EFT, the critical values are of order the cutoff. The authors eventually note that an exact higher-order completion can change the global behavior, which is correct and important. This qualification should be moved to the beginning of the global discussion. Global statements about a truncated EFT field redefinition are otherwise outside the controlled regime. The function space, topology, regularity, and transported boundary conditions should also be specified before claiming a global configuration-space diffeomorphism.

7. The scalar susceptibility requires an explicit regulator and measure calculation.

Equation (3.17) is not the complete first-order result for the flat-measure generator in equation (3.15) with a generic UV regulator. In addition to the two displayed tadpole terms, Wick contraction of $K\psi(z)$ with one of the three fields at the same interaction vertex gives

$$-6 \frac{a}{M^2} \hbar^2 \delta^{(d)}(0) \int d^d z \Delta(x, z) \Delta(z, y).$$

This term vanishes in dimensional regularization as a scaleless contact integral, but not for a generic cutoff. The Jacobian of the transformed measure produces the compensating contact term. Thus equation (3.17) is valid only after an explicit regularization convention, or after the appropriate measure factor has been included.

The authors should show this derivation and state consistently whether Bob’s generator uses a flat coordinate measure or the transformed invariant measure. The cancellation discussed in

Section 4.2 should credit both the nonlinear source and the measure; in dimensional regularization one may then explain why the ultralocal measure contribution is invisible. This is a useful opportunity to sharpen the manuscript’s central message that an action alone does not specify an off-shell quantum representation.

8. The $f(R)$ conclusion is presently conditional and largely definitional; a nontrivial test is needed.

In equation (8.45), integrating the genuine multiplier λ does impose $R - \chi = 0$, subject to the stated density convention and to factorization of the full parent measure into the chosen Fourier measure for λ and a λ -independent remainder. A generic $\sqrt{\det G_\lambda[g, \chi, \lambda]}$ would instead be Fourier transformed and would not yield a pure delta functional. Once the source is coupled only to g and the reduced measure is required to equal the direct metric measure, equation (8.46) follows. This is an exact endpoint identity, but it does not establish a new dynamical equality between two independently defined quantum theories. Equation (8.49) follows because the parent data have been chosen to reproduce the endpoint generator.

The manuscript should therefore temper the abstract and conclusion. It can correctly say that the multiplier representation is an endpoint-adapted lift of the metric path integral, while the freely integrated two-field action $\widehat{S}[g, \chi]$ is a different quantum construction unless an exact path-integral identity, including its determinant and measure, is supplied. It should not present this conditional construction as a general proof that all auxiliary-field reformulations quantize equivalently.

To establish journal-level significance, I strongly encourage an explicit calculation for a definite model, for example $f(R) = R + \alpha R^2$. The authors could compare the one-loop VDW Hessians, gauge and ghost sectors, normal determinants, and counterterm matching in the metric, multiplier-parent, Jordan, and Einstein descriptions. Such a calculation would make clear which previously reported off-shell discrepancies are ordinary noncovariant effective-action effects, which arise from different measures, and which compare genuinely different integrations. Without such a calculation, the main application is a useful bookkeeping identity but does not yet demonstrate the physical reach claimed for the framework.

9. Novelty and the relation to existing literature must be made more precise.

VDW covariance, the covariant geodesic source, the invariant measure, and gauge-condition independence are established ingredients. Likewise, the on-shell equivalence theorem, the dependence of correlators on the sourced operator, and the local nature of non-invertible changes of variables are well known. The paper’s potentially distinctive result is the combined probe-relative organization of these ingredients with parent reductions and branch gluing. That contribution should be stated plainly and separated from the reviewed material.

In particular, the paper should engage much more directly with Falls, *Equivalence of effective actions*, which analyzes how different source choices give effective actions related by an implicit mean-field transformation. The authors should explain whether the VDW criterion is a stronger requirement, a geometric realization of that correspondence, or a criterion for a different object. The discussion should also explain precisely how the proposed comparison changes the interpretation of the one-loop results of Kamenshchik–Steinwachs, Ruf–Steinwachs, Ohta, Falls–Herrero-Valea, and Finn–Karamitsos–Pilaftsis. At present these works are cited, but their conclusions are not compared equation by equation with the quantum objects defined here.

Finally, the relation $\simeq_N$ in equation (7.3) needs a sharper definition. The allowed finite parameter maps, the operator basis, composite operator renormalization, regulator dependence, possible anomalous Jacobians, and the meaning of equality on a field-space domain should be specified. Otherwise the relation is too flexible to serve as a practical test. The main statements would be clearer as propositions with explicit hypotheses and conclusions rather than as a sequence of broadly worded definitions.

Minor comments

1. In Section 2.1, $\chi = R$ follows from the λ equation in equation (2.4), not from the χ equation. The wording before equation (2.6) interchanges which variable is being eliminated. The subsequent description leading to equations (2.7)–(2.9) should be corrected consistently.
2. The paper often calls response functions and correlators off-shell observables. This phrase can conflate three distinct notions: dependence on off-shell background fields, Green functions before LSZ amputation, and physical real-time responses evaluated in a state. The terminology should be made more precise.
3. Equation (5.6) says that the covariant derivative tower contains all information about the theory. This requires analyticity and appropriate global data and is too strong without qualification.
4. The response tensor in equation (5.39) has the wrong index variance: differentiating a scalar with respect to the covector source J_A produces contravariant indices. The corresponding probe pairing should be adjusted.
5. Section 6.1 should say that the manuscript studies two representative classical probe classes. The claim that the useful possibilities reduce to only $\mathfrak{D}_{\text{kin}}^{\text{cl}}$ and $\mathfrak{D}_{\text{resp}}^{\text{cl}}$ is not justified.
6. The assertion following equations (7.13)–(7.14) that a similarity transformation preserves zero modes and propagating content requires the field redefinition to preserve operator domains and boundary conditions. For unbounded functional operators, algebraic similarity alone is not enough.
7. Equations (8.9)–(8.11) define exact pullbacks. The $O(\Lambda^{-4})$ remainder in equation (8.20) is therefore unnecessary unless the action, metric, connection, and field map have all first been truncated. The exact and truncated calculations should not be mixed.
8. In equation (8.16), the notation $\Gamma_i^{(1)}$ usually denotes the one-loop correction alone, whereas the displayed expression includes the classical action. A notation such as $\Gamma_i = S_i + \hbar\Gamma_{i,1} + O(\hbar^2)$ would be clearer.
9. The coefficient a/M^2 in the d -dimensional scalar warm-up is dimensionally appropriate with dimensionless a only in four dimensions. Either specialize to $d = 4$ or assign the correct dimension to a .

10. The absolute determinant in equations (7.25), (7.49), and related Lorentzian functional identities requires an orientation and contour prescription. The authors should say whether the argument is fundamentally Euclidean and then analytically continued.
11. The abstract should state explicitly that the claimed $f(R)$ quantum equality is conditional on a delta-constrained, endpoint-adapted source and measure. This is essential to the result, not a technical detail.
12. The manuscript is repetitive. The Introduction, both Lessons subsections, the formal definitions, the applications, and the Conclusions restate much of the same taxonomy. A table listing each probe class, the theory data it tests, and the corresponding equivalence condition would allow a substantial shortening.
13. Several small language corrections are needed: the relation between their susceptibility should read the relation between their susceptibilities; Alice and Bob state of affairs should be rewritten; the source coupling is also the pullbacks should be singular; and Eqs should be followed by a period and nonbreaking space where appropriate.

Final assessment

The manuscript has the ingredients of a useful conceptual paper, and its emphasis on comparing the same physical probe rather than the same-looking field component is valuable. The direct field-redefinition discussion and the scalar Hessian example are largely convincing. The principal parent-mediated results, however, are not established by the present definitions: the constrained integral is being conflated with ordinary auxiliary-field quantization, the ambient density is restricted to the section non-invariantly, and the endpoint correspondence is missing from the classical definition. Correcting these points, narrowing the real-time claims, and adding a concrete one-loop gravitational application would make the work substantially stronger. I would be willing to reassess a revised version that addresses these issues in full.