The Projection Operator Moves

Why John Klauder Is the Billy G.O.A.T. on This Mountain

Constraints, Reduced Kernels, BRST, and the Geometry of Physical Ranges

View from the Mountain notes from a whiteboard

Series:	A View from the Mountain, Research Note No. 3
Author:	J. Scott Little
Affiliation:	Independent Researcher
Version:	1.0.1, copyright edition
Date:	19 August 2026

Research posture and claim boundary

This note combines published mathematics, historical memory, present reconstruction, and a modern geometric reading. The labels are part of the argument. The recovered 2005 finite model proves that a physical projection can move. The lost Yang–Mills calculation is not treated as recovered fact. Its exact one-loop operator and its off-diagonal witness remain open.

Abstract

The projection-operator method for quantum constraints begins with an unusually plain act: construct a positive operator from the constraints and retain a small spectral band near zero. The physical Hilbert space is the range of that spectral projection. This note follows the consequences of taking that range seriously. Reduced reproducing kernels show how a physical space can survive a singular limiting procedure, including irregular constraints whose sectors scale at different rates. A finite anomalous model with John Klauder then provides a certified result: two classically equivalent constraint sets can lead, after quantization, to nearby but genuinely different rank-one physical ranges. The projection operator moves.

That result became the retrospective key to an unfinished field-theory calculation. In 2006–2007, alongside Wayne Bomstad and Sergei Shabanov's work on interacting QED with scalar and spinor matter, a noncovariant-gauge Yang–Mills calculation produced one-loop corrections to a squared-constraint operator. The notebooks are lost, so the exact correction is not asserted here. Instead, projector perturbation theory isolates the question that the calculation would now have to answer: does the correction have a nonzero off-diagonal block between the physical range and its complement? A window renormalization may move eigenvalues without moving the range; an off-diagonal term rotates the range itself.

The same distinction sharpens the comparison with BRST. Under closed-range Hodge hypotheses, ghost-number-zero BRST cohomology is represented by harmonic states and can coincide with the zero spectral projector of the BRST Laplacian. Outside that regime, a nonzero spectral window, continuous spectrum, second-class components, or gluing data matter. The modern picture is a family of physical Hilbert spaces over regulators, windows, observers, and times, equipped when possible with the Kato connection $P\mathrm{d}$. Motion is then not a defect to hide. It is geometry to record. Klauder's enduring achievement was to make the physicalizing operation visible before the modern language was ready for it.

Evidence labels used throughout

Recorded means supported by a surviving publication or document. **Remembered** means autobiographical memory used for chronology and motivation, not for a mathematical priority claim. **Reconstructed** means an argument rederived here from stated assumptions. **Modern reading** means a later geometric interpretation. **Open** marks a claim that the surviving evidence does not settle.

Reconstruction notice

The missing notebooks are not evidence. Reproduced mathematics can certify where a line of thought had reached only when it agrees with surviving work. Where that agreement stops, this note stops. The history is therefore neither an archive dump nor a polished legend; it is a ledger of what can still be known.

The result is certifiable or it is not.
The progression of thought remains invaluable.

Contents

Scope firewall

This note does not claim to recover the lost Yang–Mills calculation, solve interacting Yang–Mills theory, replace BRST, or construct a complete global field-theoretic functor. It proves and reconstructs narrower statements: how reduced kernels survive, how a finite physical projector moves, when BRST and a spectral projector agree, and what geometric data a moving family carries.

1 The notebooks are gone, but the lesson survived

1.1 Old mathematics as a trail marker

There is a peculiar kind of historical evidence available to a working physicist. A notebook can disappear while the habits learned in it remain. Years later a calculation can be reconstructed, and a surviving paper can certify that certain ideas were already in hand even when the route between them is gone. This is weaker than an archive and stronger than nostalgia. Used carefully, it can tell us where we were.

The organizing question of this note is simple to state: if the constraints change, does the projection onto physical states merely acquire a new normalization, or does its range rotate inside a larger Hilbert space? At the time, the second possibility looked like a failure. Why should the operator that selects the physical states need a quantum correction? Why should it move? From the mountain, the premise has reversed. If the physical condition depends on scale, regulator, representation, time, or quantum anomaly, why would its carrier remain rigid?

Mountain memory: Shabanov's question

REMEMBERED. Sergei Shabanov asked, "What does this say about quantum field theory?" My answer was: "Nothing — yet." The word *yet* was the honest part. A finite model could expose the mechanism without pretending that the field theory had already been controlled.

The story begins below field theory, with a spectral band and a reproducing kernel. It reaches field theory only after the finite argument has taught us what must be measured. This order matters. It prevents a remembered loop calculation from doing work that only a surviving theorem can do.

1.2 Florida, seed work, and the mountain view

For years I regarded the early constraint papers as exercises. They were finite, explicit, and therefore easy to rank below the grand problems. In a modern light they look different. The reduced kernel already asks how a physical carrier is extracted from a limiting family. The anomalous model already asks how that carrier changes under deformation. The irregular model already refuses a single universal scaling. Those are not the entire modern theory, but they are unmistakably its seeds.

There is a personal correction here as well. One can absorb a damaging idea that serious physics belongs only to a handful of institutional addresses. That idea is false. Florida was a very good school. John Klauder, Wayne Bomstad, Sergei Shabanov, and the problems themselves demanded real physics. The work did not become meaningful retroactively because a modern vocabulary can now describe it. The modern vocabulary lets me see more clearly that the questions were already real.

Thesis in one sentence

The projection operator moves when the quantum physical subspace changes its position in the ambient Hilbert space; a spectral shift is not enough, but an off-diagonal coupling between the old physical range and its complement is.

2 Why write a reconstruction in the age of instant answers?

2.1 The result is certifiable or it is not

The loss of the notebooks creates a temptation in two opposite directions. One can say nothing, because the original route cannot be reproduced word for word. Or one can smooth the memory into a finished story and let confidence stand where evidence should have stood. Neither response is satisfactory.

The better rule is severe and simple: the result is certifiable or it is not. A surviving publication can certify a theorem. A new calculation can certify that a mechanism works under newly stated assumptions. A remembered conversation can explain motivation when it is labeled as memory. None of these can silently certify the others.

That rule does not make the progression of thought disposable. Quite the opposite. The progression tells us why a question was visible, which wrong interpretation blocked it, and which later idea made it legible. A scientific result stripped of that progression may remain correct, but it can lose the information that helps another person recognize the next problem.

This series is therefore not an attempt to turn every old calculation into a priority claim. It is an attempt to preserve the inquiry itself while allowing formal proof to decide which mathematical statements survive.

2.2 Why this matters more with language models

We now possess tools that are remarkably unwilling to leave a blank space blank. Ask for a missing derivation and a plausible derivation may appear. Ask what an old result "must have meant" and a coherent interpretation can be supplied before the historical uncertainty has had time to become uncomfortable.

That ability is useful. It is also exactly why reconstruction notices are needed. Fluency does not restore a notebook. A complete-looking proof is not evidence that the same proof existed twenty years earlier. A plausible operator is not the regulated operator that was actually calculated.

The appropriate use of the tool is narrower and stronger. It can help stage the alternatives, reproduce mathematics from explicit assumptions, test whether a remembered phenomenon is possible, and expose which missing datum would decide among several interpretations. It can help turn memory into an audit. It cannot turn memory into an archive.

The labels Recorded, Remembered, Reconstructed, Modern reading, and Open are therefore not decorative. They are interfaces between different kinds of warrant. Removing them would make the prose smoother and the history less true.

2.3 Insight is not an optional ornament

Scientific journals increasingly contain results whose authors candidly say that they supplied little or no conceptual insight into the discovery. That is a choice. It may still produce correct mathematics. But it is not the choice I want for this series.

I like the new tools. I also trust the odd and imperfect instrument between my ears as the source of long-term conjecture, discomfort, and direction. It is the instrument that notices when two calculations smell alike before there is a theorem connecting them. It is the instrument that remains bothered by a projection operator moving when the formal calculation says everything has closed.

Intuition is not certification. It can be wrong for years. But removing it from the account makes the eventual theorem appear to have arrived without a question. The history of a problem is not only the history of correct statements. It is the history of learning what must be proved.

Mountain memory: The purpose of the View from the Mountain series

The notes preserve the trail: early seeds, wrong turns, surviving calculations, later reinterpretations, and the exact point at which the present evidence stops. The view from the mountain is wider than the view on the trail, but it must not redraw the trail merely because the wider view is prettier.

2.4 Old mathematics as a witness to location

Reproducing an old calculation today does not prove every detail of when it was first obtained. It does something more limited. When the reconstruction matches a surviving paper, dissertation, or dated correspondence, it can certify the level of mathematical control that had been reached. The old publication becomes a trail marker.

The anomalous model is such a marker. Its rank-one projector moves by a computable amount. That fact is not a present invention. What is modern is the decision to measure the motion with Grassmannian language and to recognize it as the finite prototype of the later question.

The reduced-kernel work is another marker. It establishes that a physical Hilbert space can survive a singular projection limit after the correct renormalization, even when the original projection tends to zero. The modern language of varying carriers makes the lesson easier to place, but the kernel mechanism itself was already there.

The lost Yang–Mills calculation is not such a marker. It is a remembered branch leaving the documented trail. This note can identify the operator block that a successful reconstruction must compute. It cannot report the value of that block until the calculation has been performed again.

That difference is the moral discipline of the series. It also gives the notes their value. A reader sees not only an answer, but the boundary between what was proved, what can now be reproved, and what remains on the mountain.

3 Constraints in one spectral window

3.1 From many constraints to one positive operator

Let $\{\Phi_\alpha\}$ be quantum constraint operators on an auxiliary Hilbert space $\mathcal{H}$, and let $G^{\alpha\beta}$ be a positive matrix or positive kernel on the constraint labels. On a common invariant form domain define

$$C = \sum_{\alpha,\beta} \Phi_\alpha^\dagger G^{\alpha\beta} \Phi_\beta, \qquad C \geq 0. \tag{1}$$

For $\delta > 0$, the projection-operator prescription is

$$P_\delta = \chi_{[0,\delta^2]}(C), \qquad \mathcal{H}_{\mathrm{phys},\delta} = \mathrm{Ran}\, P_\delta. \tag{2}$$

The spectral window replaces the formal demand that every constraint annihilate every physical vector. It is meaningful for first-class constraints, second-class constraints, anomalous mixtures, and cases where zero belongs to continuous spectrum. The price of this breadth is that the window is data, not punctuation.

What the spectral band does

Equation (2) separates three operations that are often blended: constructing the quantum constraints, forming a positive diagnostic C, and choosing which part of its spectrum counts as physical. Renormalizing C, changing δ, and rotating $\mathrm{Ran}\, P_\delta$ are therefore logically distinct events.

3.2 One spectral direction through $N^2 - 1$ color directions

The squared-constraint construction does something more important than make a positive operator. It compresses a multidirectional constraint problem into one spectral direction. For a compact gauge group with Lie algebra $\mathfrak{g}$, write the constraint family as a vector-valued operator

$$\mathbf{\Phi} : \mathcal{H} \longrightarrow \mathcal{H} \otimes \mathcal{K}_{\mathrm{con}}. \tag{3}$$

For $SU(N)$ there are $N^2 - 1$ finite-dimensional color directions before the spatial labels of a field theory are even counted. The componentwise Dirac condition reads

$$\Phi_a \psi = 0, \qquad a = 1, \ldots, N^2 - 1, \tag{4}$$

and in Yang–Mills theory it becomes

$$G_a(x)\psi = 0 \quad \text{for every color } a \text{ and admitted spatial smearing.} \tag{5}$$

This is a forest of directions. The positive contraction

$$C = \mathbf{\Phi}^\dagger G \mathbf{\Phi} \tag{6}$$

turns the forest into the single nonnegative spectral axis of C.

On the common form domain,

$$\langle \psi, C\psi \rangle = \|G^{1/2}\mathbf{\Phi}\psi\|^2. \tag{7}$$

If G is strictly positive on the admitted constraint-label space, then

$$\operatorname{Ker} C = \operatorname{Ker} \mathbf{\Phi} = \bigcap_a \operatorname{Ker} \Phi_a. \tag{8}$$

Thus one does not have to form $N^2 - 1$ separate physical projectors and then discover how to intersect them. One studies one positive operator and one window.

This is what I meant by looking in one direction. It is not one surviving gauge-generator direction. It is one *spectral* direction carrying the combined norm of all constraint violations. In a local field theory the compression is more dramatic: the color and spatially smeared directions are all collected into a single regulated quadratic form.

Constraint compression

Let $\mathbf{\Phi}$ be closed and densely defined, and let the positive quadratic form

$$q_C[\psi] = \|G^{1/2}\mathbf{\Phi}\psi\|^2 \tag{9}$$

be closed. The associated nonnegative self-adjoint operator C satisfies $\operatorname{Ker} C = \operatorname{Ker} \mathbf{\Phi}$. Consequently the exact common constraint kernel is the zero spectral space of one operator, while approximate satisfaction of all constraints is represented by one low spectral band.

The domain qualification matters. A formal sum of unbounded squares is not automatically a self-adjoint operator, and positivity alone does not repair an undefined product. Klauder's architecture is most naturally understood at the level of closed positive forms. Once the form is controlled, the spectral theorem supplies the projector.

3.3 Where the representations went

The compression does not erase representation theory. Suppose a compact group acts on

$$\mathcal{H} \cong \bigoplus_R \mathcal{M}_R \otimes V_R, \tag{10}$$

where V_R is irreducible and $\mathcal{M}_R$ is its multiplicity space. For an invariant metric and orthonormal generators T_a, the quadratic Casimir acts as

$$\sum_a T_a^{(R)} T_a^{(R)} = c_2(R)\mathbf{1}_R. \tag{11}$$

The representation directions are reorganized along the single Casimir spectrum $R \mapsto c_2(R)$. The trivial representation lies at zero, while nontrivial representations occur at positive Casimir values.

For a composite system the total generator is

$$T_a^{\text{tot}} = \sum_j T_a^{(j)}, \tag{12}$$

and its square expands as

$$C_{\text{tot}} = \sum_j \sum_a T_a^{(j)} T_a^{(j)} + 2 \sum_{j<k} \sum_a T_a^{(j)} T_a^{(k)}. \tag{13}$$

The cross terms remember how the separate representations fuse. The zero-Casimir sector selects the singlets. The representation geometry has not gone away; it has been concentrated into eigenvalues, multiplicities, degeneracies, and embeddings of spectral sectors.

This was especially valuable in the problems surrounding interacting matter. Scalar matter, spinor matter, and gauge fields carry different pieces of the total gauge action. The squared constraint asks whether their total content lands in the singlet band without requiring every representation direction to be followed independently through every intermediate expression.

Audit: what one spectral direction does not prove

The equality $\operatorname{Ker} C = \bigcap_a \operatorname{Ker} \Phi_a$ does not establish anomaly-free closure of the component algebra, regulator independence, or preservation of domains. The non-Abelian information survives inside C even though it has been spectrally compressed. One positive direction simplifies the question; it does not make the gauge algebra Abelian.

3.4 Why this mattered for the loop calculation

Suppose the constraint vector receives a perturbative correction,

$$\mathbf{\Phi}(g) = \mathbf{\Phi}_0 + g\mathbf{\Phi}_1 + O(g^2). \tag{14}$$

Then

$$C(g) = C_0 + gC_1 + g^2C_2 + \cdots, \tag{15}$$
$$C_0 = \mathbf{\Phi}_0^\dagger G \mathbf{\Phi}_0, \tag{16}$$
$$C_1 = \mathbf{\Phi}_1^\dagger G \mathbf{\Phi}_0 + \mathbf{\Phi}_0^\dagger G \mathbf{\Phi}_1. \tag{17}$$

For $\psi_0 \in \operatorname{Ker} \mathbf{\Phi}_0$,

$$\langle \psi_0, C_1\psi_0 \rangle = 0, \tag{18}$$

but

$$C_1\psi_0 = \mathbf{\Phi}_0^\dagger G \mathbf{\Phi}_1 \psi_0 \tag{19}$$

need not vanish. The diagonal first-order energy shift may therefore be zero while the off-diagonal mixing is nonzero. The physical range can move at first order even when the small eigenvalue does not lift until second order.

That is a crucial feature of the old problem. Watching only the lowest eigenvalue would miss the first geometric effect. The squared constraint made the many directions manageable, but it also made it necessary to distinguish spectral displacement from motion of the spectral subspace.

If zero is an isolated eigenvalue and $0 < \delta^2$ lies below the next spectral point, then $P_\delta = \chi_{\{0\}}(C)$. If zero is not an eigenvalue, the projections may converge strongly to zero as $\delta \downarrow 0$. That does not automatically mean that the physical theory disappears. It means the ambient Hilbert norm and the limiting physical norm need not be the same. The reproducing-kernel construction makes this explicit.

3.5 What it means for a projector to move

Suppose $P(s)$ is a differentiable family of orthogonal projections. From $P^2 = P$ one obtains

$$P\dot{P}P = 0, \qquad \dot{P} = P\dot{P}(1-P) + (1-P)\dot{P}P. \tag{20}$$

Thus the tangent to the Grassmannian is purely off-diagonal relative to $\mathcal{H} = \operatorname{Ran} P \oplus \operatorname{Ker} P$. A change within $P\mathcal{H}$ is a choice of frame; a nonzero $(1-P)\dot{P}P$ is motion of the subspace itself.

For finite-rank projections the Hilbert–Schmidt speed is

$$\|\dot{P}\|_{\mathrm{HS}}^2 = 2\|(1-P)\dot{P}P\|_{\mathrm{HS}}^2. \tag{21}$$

This elementary identity will later turn a qualitative phrase into a measured fact in the 2005 anomalous model.

Audit: small eigenvalues do not by themselves imply motion

If $C(s) = a(s)C(0)$ with $a(s) > 0$, its eigenvalues and the appropriate spectral window may change while every spectral subspace stays fixed. Motion requires eigenvector mixing, rank change, or a change of ambient identification. The first-order witness is off-diagonal.

4 Reduced reproducing kernels and fixed physical points

4.1 The projected coherent-state kernel

Let $\{|z\rangle : z \in X\}$ be coherent states with resolution of the identity

$$\int_X |z\rangle\langle z| \, \mathrm{d}\mu(z) = \mathbf{1}_{\mathcal{H}}. \tag{22}$$

The projected kernel

$$K_\delta(z'', z') = \langle z''|P_\delta|z'\rangle \tag{23}$$

is positive definite. Before any singular limiting operation, orthogonality of P_δ gives the exact reproduction identity

$$\int_X K_\delta(z'', z) K_\delta(z, z') \, \mathrm{d}\mu(z) = K_\delta(z'', z'). \tag{24}$$

The range of P_δ is therefore represented as a reproducing-kernel Hilbert space. This is not merely a convenient coordinate system. It is a way to keep the physical inner product visible while the ambient spectral projection becomes singular.

4.2 The reduction theorem

Assume that positive scale factors $S(\delta)$ can be chosen so that

$$\widetilde{K}(z'', z') = lim_{\delta\downarrow 0} S(\delta) K_\delta(z'', z') \tag{25}$$

exists pointwise and is not identically zero.

Reduced-kernel positivity

Every pointwise limit (25) of rescaled positive kernels is a positive kernel. It therefore defines a reproducing-kernel Hilbert space after quotienting null vectors and completing.

Proof. For any finite set $z_1, \ldots, z_n$ and coefficients $c_1, \ldots, c_n$,

$$\sum_{i,j} \bar{c}_i c_j \widetilde{K}(z_i, z_j) = \lim_{\delta\downarrow 0} S(\delta) \sum_{i,j} \bar{c}_i c_j K_\delta(z_i, z_j) \geq 0. \tag{26}$$

The Moore–Aronszajn construction supplies the Hilbert space. □

The point that requires care is what has *not* been proved. After a singular rescaling, $\widetilde{K}$ need not satisfy (24) with the original measure μ. The limiting RKHS has its own native reproduction law. Positivity survives directly; idempotence with respect to the old ambient resolution may not.

4.3 Irregular constraints require more than one magnifying glass

Suppose the small-window space decomposes into asymptotically orthogonal sectors,

$$K_\delta = K_\delta^{(1)} + \cdots + K_\delta^{(r)} + o(1), \tag{27}$$

and sector a vanishes like $s_a(\delta)^{-1}$. A single scale factor can erase all but the dominant sector. The correct reduction may instead be componentwise:

$$\widetilde{K} = \bigoplus_{a=1}^{r} \lim_{\delta \downarrow 0} s_a(\delta) K_\delta^{(a)}. \tag{28}$$

This is the lesson of the highly irregular examples developed with Klauder [7]. The physical limit is not always one fixed point of one scalar renormalization. It may be a direct sum of fixed points visible at different rates.

Audit: universality of the window

The method is universal in the sense that one spectral prescription can treat many constraint types. It is not universal in the stronger sense that one power of δ must normalize every physical sector. Irregularity is data about the limiting geometry.

5 A certified motion: the 2005 anomalous model

5.1 Classical equivalence, quantum inequivalence

The surviving paper *Elementary Model of Constraint Quantization with an Anomaly* supplies the decisive historical anchor [6]. Classically, let j_i be angular-momentum-type constraints and let

$$l_i = f(q) j_i \tag{29}$$

with $f(q) > 0$. The two families have the same classical zero set. Their quantizations need not define the same constraint algebra. In the model, the J_i remain first class, while the symmetrically ordered

$$L_i = AJ_i + J_i A \tag{30}$$

form a partially second-class family. The difference is not in the classical constraint surface but in the quantum physicalizing operator.

For the unmodified system the low spectral band is rank one:

$$P_J = |O_J\rangle\langle O_J|, \qquad |O_J\rangle = \frac{|200\rangle + |020\rangle + |002\rangle}{\sqrt{3}}. \tag{31}$$

For small deformation parameter β, the modified physical vector is

$$|O_L(\beta)\rangle = M(\beta)\Big(d(\beta)|200\rangle + d'(\beta)|020\rangle + |002\rangle\Big), \tag{32}$$

with

$$d(\beta) = 1 - 2\beta + O(\beta^2), \qquad d'(\beta) = 1 - \beta + O(\beta^2). \tag{33}$$

Normalization gives

$$|O_L(\beta)\rangle = |O_J\rangle + \frac{\beta}{\sqrt{3}}\Big(-|200\rangle + |002\rangle\Big) + O(\beta^2). \tag{34}$$

The derivative is orthogonal to $|O_J\rangle$. It is therefore not a phase or a change of basis inside a fixed one-dimensional range.

Finite anomalous motion

Let $P_L(\beta) = |O_L(\beta)\rangle\langle O_L(\beta)|$. Then

$$\dot{P}_L(0) \neq 0, \qquad g_{\beta\beta}(0) = \langle \dot{O}_L|\dot{O}_L\rangle - |\langle O_L|\dot{O}_L\rangle|^2 = \frac{2}{3}, \tag{35}$$

and

$$\|\dot{P}_L(0)\|_{\mathrm{HS}}^2 = \frac{4}{3}. \tag{36}$$

Thus the physical line moves in projective Hilbert space at nonzero speed.

This is the proof I did not know I would need later. It says exactly what the phrase "the projection operator moves" should mean. The rank remains one; the line changes position. No change of window alone can account for that fact.

5.2 The small eigenvalue and the nonzero window

In the modified model, the bottom of the master-constraint spectrum is lifted. For the regime used in the paper,

$$\lambda_{\min} \sim \frac{1}{2}\hbar^2\beta^6, \qquad \delta^2 = \beta\hbar^2, \tag{37}$$

so the low state remains inside a nonzero spectral band. This is not a technical embarrassment to be removed before interpreting the result. It is how a partially second-class quantum system retains a physical sector.

Methods that depend only on the classical zero set can be insensitive to the factor f. The projection method is deliberately sensitive to the quantum operator obtained after ordering. Its answer changes because the quantum constraint problem changed.

Mountain memory: Wayne's excitement and formal mathematics

Remembered. Wayne was excited by the reduced-kernel and anomalous calculations because formal mathematics had entered the fray at the point where the physics demanded it. The finite model was not a toy to be discarded. It was a controlled place where anomaly, range selection, and deformation could all be seen at once.

6 Wayne, Shabanov, and the road from QED to Yang–Mills

6.1 The live frontier was interacting QED

The historical order matters. Wayne Bomstad and Sergei Shabanov were working on the projection-operator treatment of QED interacting with scalar and spinor matter. In schematic canonical form, the Gauss constraint is

$$G_{\mathrm{QED}}(x) = \nabla \cdot E(x) - e\Big(j^0_\phi(x) + j^0_\psi(x)\Big). \tag{38}$$

Matter pushes on the gauge constraint, and loop corrections test whether the regulated physical band can be controlled without first reducing the classical phase space.

By 2006, before I left Florida and got married, Wayne's loop calculations were starting to go sideways. He continued. During brief quarterly visits — Thursday-to-Sunday whirlwind tours to meet with John and report back — Wayne would sit down with me and talk through his discussions. I pointed out techniques. He cautioned that he wanted to do it his own way. I told him I had an adjacent problem in mind and would keep him posted.

6.2 The Thursday-to-Sunday visits

Those visits deserve more than a chronological line because they show how the problem actually developed. I had left Gainesville, married, and entered a different rhythm of life. The work no longer happened simply because I could walk down a hall. It had to be carried across distance in short, concentrated bursts.

I would return for a long weekend, usually arriving on Thursday and leaving on Sunday. The visit had several purposes. I needed to meet with John, show him what I had done, and report back. Wayne had his own calculations and his own questions. A few days therefore had to contain the conversations that once could have occurred gradually over weeks.

Wayne would sit down and talk through the QED calculation. He was not merely asking me to debug algebra. He was explaining where the calculation had begun to resist the conceptual story that should have organized it. I would suggest techniques or point to a way of separating terms. Sometimes a second pair of eyes can see that two apparent problems are the same calculation in different clothing.

But Wayne also made a boundary clear: he wanted to do the problem in his own way. I respected that. Independence matters in research, especially when the calculation is close enough to another person's work that advice can become ownership without anyone intending it.

My answer was that I had an adjacent problem in mind. I would work beside the QED problem rather than inside it, and I would keep him posted. This is the historically correct doorway through which Yang–Mills enters the story. It was not a replacement for Wayne's calculation and not an attempt to take it over. It was a neighboring test of the same projection-operator architecture in a theory where the gauge field participates in its own constraint.

Mountain memory: Research beside a friend

The exchange was collegial and careful. Wayne wanted responsibility for his route through QED. I wanted to help without swallowing the problem. The adjacent Yang–Mills calculation let both things remain true.

6.3 Enter Yang–Mills

For Yang–Mills, the Gauss constraint is already self-interacting:

$$G^a_{\mathrm{YM}}(x) = (D_i E_i)^a(x) = \partial_i E^a_i(x) + g f^{abc} A^b_i(x) E^c_i(x). \tag{39}$$

The adjacent problem was therefore adjacent in exactly the dangerous way. The gauge field appears in its own constraint. The master operator built from G^a_{YM} probes interactions even before matter is added.

The handoff in one line

In QED, matter pushes on the gauge constraint. In Yang–Mills, the gauge field pushes on its own constraint.

For Christmas the preceding year I had received Bassetto, Nardelli, and Soldati's *Yang–Mills Theories in Algebraic Non-Covariant Gauges* [1]. I used its treatment of algebraic noncovariant gauges to look at the squared constraints and their quantum corrections through one-loop order. The noncovariant choice was deliberate. I did not want manifest covariance to mask the real canonical difficulties.

6.4 A Christmas book and an honest gauge

Books sometimes enter a research history as if they had been selected after the theorem was known. This one entered as a Christmas gift. Its subject made it immediately useful because the adjacent problem was already taking shape. Algebraic noncovariant gauges kept the canonical constraint structure near the surface of the calculation.

My reason for using such a gauge was not hostility to covariance. Covariance is powerful, and manifestly covariant formalisms often organize perturbation theory beautifully. But the question I wanted to see was precisely the one a smooth formalism might encourage me to pass over: what had happened to the operator selecting the physical states?

In a noncovariant canonical presentation, Gauss law does not disappear behind notation. One can see which variables are constrained, which terms couple the constraint sectors, and where ordering or loop corrections enter the squared operator. The gauge therefore served an epistemic purpose. It made the difficult part harder to ignore.

This choice did not eliminate ghosts, singular denominators, residual gauge issues, or renormalization questions. It simply placed the projection problem where I could look at it directly. I wanted the calculation to fail honestly if it was going to fail.

Audit: the exact gauge choice

The surviving memory fixes the use of an algebraic noncovariant gauge, but not whether the completed notebook calculation used the temporal, axial, light-cone, or a more general member of that family. The exact gauge is Open. This note does not select one retroactively.

The calculation was completed, but I did not understand its conclusion. The squared-constraint projection seemed to require quantum corrections. The projection itself appeared to move. I considered looking at the renormalization group; it looked promising and difficult. Because I did not yet have the language to separate a moving eigenvalue from a moving range, I thought I had failed to close the calculation. In retrospect, the right question was already there.

6.5 A calculation can close before its interpretation does

The phrase "I did not close it" needs correction. There are several forms of closure in a constrained quantum calculation. The constraint algebra may close. The perturbative equations may close. The counterterm list may close. The chosen spectral band may remain separated. And yet the physical range may fail to be the same range with which the calculation began.

At the time I treated the last event as evidence against the others. If the physical projection needed a correction, I assumed I must have mishandled the constraint algebra. The 2005 finite model had already shown otherwise, but I had not extracted its geometric lesson. A closed quantum constraint problem can define a moving physical subspace.

The renormalization group appeared promising because both the constraint operator and its window could depend on scale. What I did not yet see was that the comparison between scales was an additional object. It was not enough to compute C_Λ at every cutoff. One had to say how $P_\Lambda \mathcal{H}_\Lambda$ at one scale was to be identified with the physical space at another.

That question was much larger than the dissertation calculation I was trying to finish. Recognizing that it was larger would eventually matter. At the time it mostly felt like not understanding what I had done.

6.6 Linearized gravity as a neighboring success

Bomstad and Klauder's linearized-gravity analysis showed what the projection method could already do in a field setting [2]. It treated an infinite family of constraints, including a mixture with second-class behavior, without inserting a classical gauge-fixing stage. The physical transverse-traceless sector emerged from the spectral construction. That result did not solve interacting QED or Yang–Mills. It established that the projection method could carry canonical field constraints without collapsing into a finite-dimensional trick.

7 The one-loop question, reconstructed honestly

7.1 Regulated master operators

Let Λ summarize the ultraviolet regulator, finite volume, smearing, and gauge parameters needed to define a self-adjoint master operator. A renormalized one-loop calculation has the schematic form

$$C_{\Lambda,R} = C_{\Lambda,0} + \hbar C_{\Lambda,1} + O(\hbar^2), \qquad P_{\Lambda,R} = \chi_{[0,\delta_R^2]}(C_{\Lambda,R}). \tag{40}$$

Suppose a contour Γ separates the chosen low spectral cluster from the rest. With $R_0(z) = (z - C_{\Lambda,0})^{-1}$,

$$P_{\Lambda,R} = \frac{1}{2\pi i} \oint_\Gamma (z - C_{\Lambda,R})^{-1}\, \mathrm{d}z, \tag{41}$$

and hence

$$P_{\Lambda,1} = \frac{1}{2\pi i} \oint_\Gamma R_0(z) C_{\Lambda,1} R_0(z)\, \mathrm{d}z. \tag{42}$$

The off-diagonal block

$$W_{\Lambda,1} = (1 - P_{\Lambda,0}) C_{\Lambda,1} P_{\Lambda,0} \tag{43}$$

is the decisive first-order witness.

First-order range-motion criterion

Assume the low spectral cluster of C_0 is isolated by a positive gap. Then the first-order spectral projector is stationary if and only if the off-diagonal blocks of C_1 between the cluster and its complement vanish. Equivalently,

$$(1 - P_0) P_1 P_0 = 0 \quad \Longleftrightarrow \quad (1 - P_0) C_1 P_0 = 0. \tag{44}$$

For an eigenvector $|i\rangle$ in the low cluster and $|a\rangle$ outside it,

$$\langle a|P_1|i\rangle = -\frac{\langle a|C_1|i\rangle}{\lambda_a - \lambda_i}. \tag{45}$$

This criterion explains the old confusion. A counterterm proportional to C_0, or a rescaling absorbed into δ_R, changes the spectral calibration but not the physical range. A term C_1 that couples old physical states to old unphysical states rotates the range. The latter is not cured by renaming the window.

7.2 What is known and what remains open

The notebook result can now be divided into claims of different strength.

Label	Claim
Remembered	A noncovariant-gauge Yang–Mills squared-constraint calculation was carried through one loop, and the physical projection appeared to require a quantum correction.
Reconstructed	Any such calculation can be audited by separating window renormalization from the off-diagonal witness $W_{\Lambda,1}$.
Recorded	The 2005 anomalous model exhibits a nonzero moving physical range under a controlled deformation.
Open	The exact operator $C_{\Lambda,1}$ in the lost Yang–Mills calculation, the gauge member used, and whether $W_{\Lambda,1} \neq 0$ after all counterterms and limits.

The go/no-go decision for the essay does not depend on recovering the notebook. The finite model certifies the mechanism. The field calculation motivates a precise reconstruction program. If the old result is ever reproduced, it must pass the witness test (43); if it does not, the honest conclusion will be that the spectrum moved but the range did not.

8 Why not simply BRST?

8.1 What BRST genuinely accomplishes

For first-class constraints with suitable closure, BRST introduces a graded state space and a nilpotent charge

$$Q = c^a \Phi_a - \frac{1}{2} f_{ab}{}^c c^a c^b \mathcal{P}_c + \cdots , \qquad Q^2 = 0. \tag{46}$$

Physical states at ghost number zero are cohomology classes,

$$H^0(Q) = \frac{\operatorname{Ker}(Q : \mathcal{K}^0 \to \mathcal{K}^1)}{\operatorname{Ran}(Q : \mathcal{K}^{-1} \to \mathcal{K}^0)}. \tag{47}$$

This is a powerful algebraic resolution of gauge redundancy. Ghosts encode Jacobian determinants in the path integral, the homological resolution of the constraint ideal, and, in global settings, failures of a single gauge slice. None of that is denied here.

The reservation is architectural. The standard physical-state story begins with one enlarged graded complex and obtains the physical theory by a global quotient. A local family of positive physical carriers, their changes under windows and regulators, and their gluing are not the primary visible objects. Local BRST cohomology certainly exists, and BRST is not "nonlocal" in any simple sense. The point is narrower: the projection method places the positive physical range at the front of the construction.

Audit: ghosts

It is too strong to say that ghosts arise only from a conflict between local and global considerations. They have perturbative determinant and algebraic roles even on a contractible patch. Global gauge geometry explains why a single ghost-extended description may fail to descend to one global slice; it does not exhaust the origin of ghost fields.

8.2 The local-positive question

The projection method asks for $P_b \mathcal{H}_b$ at each admissible base point b. The label b may include a region, scale, regulator, time, representation, and window. The first object is then a positive Hilbert space, not an equivalence class in an indefinite or graded extension. Comparisons between these spaces can fail, be partial, or carry holonomy. That failure is visible rather than absorbed into a single quotient.

This is not a verdict that one method is right and the other wrong. It is a test of equivalence. When does the BRST cohomology admit a canonical positive harmonic representative space, and when does that space agree with Klauder's spectral range?

9 The BRST–projection bridge

9.1 Hodge representatives

Let

$$\cdots \xrightarrow{Q_{n-1}} \mathcal{K}^n \xrightarrow{Q_n} \mathcal{K}^{n+1} \xrightarrow{Q_{n+1}} \cdots, \qquad Q_n Q_{n-1} = 0, \tag{48}$$

be a Hilbert complex of closed densely defined operators. Define the BRST Laplacian at degree n by

$$\Delta_n = Q_n^\dagger Q_n + Q_{n-1} Q_{n-1}^\dagger. \tag{49}$$

BRST–Klauder equivalence under Hodge hypotheses

Assume $\operatorname{Ran} Q_{n-1}$ and $\operatorname{Ran} Q_n$ are closed, so that the Hilbert-complex Hodge decomposition holds:

$$\mathcal{K}^n = \operatorname{Ran} Q_{n-1} \oplus \operatorname{Ker} \Delta_n \oplus \operatorname{Ran} Q_n^\dagger. \tag{50}$$

Then every cohomology class in $H^n(Q)$ has a unique harmonic representative and

$$H^n(Q) \cong \operatorname{Ker} \Delta_n. \tag{51}$$

If zero is isolated in $\operatorname{Spec} \Delta_n$ by a gap $\gamma > 0$, then for $0 < \delta^2 < \gamma$,

$$\chi_{[0,\delta^2]}(\Delta_n) = \chi_{\{0\}}(\Delta_n), \tag{52}$$

so the BRST cohomology is represented exactly by a spectral projection.

The proof is the standard Hodge argument: closed ranges make exact and coexact subspaces orthogonally complemented; harmonic vectors are precisely closed vectors orthogonal to exact vectors; the spectral gap identifies the zero kernel with a finite window. The result is important because it turns an often vague comparison into a theorem with visible hypotheses.

For a Klauder master operator C_Φ, exact physical equivalence additionally requires an isometry

$$U : \operatorname{Ran} \chi_{[0,\delta^2]}(C_\Phi) \longrightarrow \operatorname{Ker} \Delta_0 \tag{53}$$

that intertwines the physical observables and dynamics. Equality of vector spaces without equality of representation data is not equivalence of physical theories.

9.2 Where exact equivalence can fail

The theorem also identifies the fault lines:

- If $\operatorname{Ran} Q$ is not closed, cohomology may be non-Hausdorff or require a reduced quotient, while zero can accumulate in the Laplacian spectrum.
- If zero lies in continuous spectrum, a nonzero Klauder window can carry states even when $\operatorname{Ker} \Delta = \{0\}$ in the ambient norm.
- Second-class or anomalous constraints may not admit the simple nilpotent BRST charge assumed above without conversion or extension.
- Even fiberwise equivalence can fail to preserve comparison maps, local gluing, time evolution, or regulator transport.

Thus "projection equals BRST" is neither universally true nor generally false. It is true on a theorem-sized domain. Beyond that domain, the nonzero-window and comparison data are exactly what must be examined.

Mountain memory: May 2007

Remembered. In May 2007, after deciding on a more limited dissertation built around collections and comparisons, I showed John a similar BRST argument together with my first time-dependent results. He was very encouraged. I told him a little about the Yang–Mills result but said I did not think it was ready for prime time. He agreed: those results could wait. If I planned to finish by the fall semester, I had to work with what was in hand.

Mountain memory: John's instruction

Remembered. "Learn what you can and move." I went home and devised a chapter-by-chapter outline. Around July, a year into marriage, John preliminarily agreed and asked to see the first three chapters before giving the full okay. The afternoons became a hard-core writing schedule. I lost roughly half the dissertation and rewrote it from notes. It was better the next time.

The memoir belongs here because the mathematical statement and the practical advice have the same structure. An unfinished global synthesis should not prevent a controlled local result from being completed. Preserve what is certified, mark what is open, and move.

9.3 What was in hand

The dissertation decision was not a rejection of the larger program. It was a statement about admissible scope. I had collections of constraint examples, comparisons between physicalization methods, a BRST bridge under controlled hypotheses, and the beginning of a time-dependent picture. Those pieces could be written, checked, and defended.

The Yang–Mills calculation was different. I could describe what I had done and show John pieces of the result, but I could not yet explain why the projection should move or which part of the correction was invariant. A calculation that cannot yet be interpreted may still be right, but it is not ready to carry a dissertation whose completion has a deadline.

John's advice separated learning from possession. I did not have to possess the entire Yang–Mills problem before completing the part I understood. I could learn what the difficult calculation was trying to teach and move with the mathematics that was already stable.

I went home and made the outline chapter by chapter. This was more than a table of contents. It was a decision about proof ownership: which results had complete arguments, which comparisons needed qualification, and which questions would be left outside the dissertation. Around July John wanted to see the first three chapters before giving the final approval. The request was both encouragement and test. An outline becomes a dissertation only when the prose and mathematics can survive contact with pages.

The afternoons became the writing schedule. By then I had been married a year, and the work had to fit a life rather than assume life would wait. Then roughly half of the dissertation disappeared. I reconstructed it from notes. The rewritten version was better.

That episode is almost too neat as a metaphor for this note, but it happened. Lost mathematics is not automatically improved by being lost. The loss is real. Yet reconstruction forces one to identify which steps were structural and which sentences merely accompanied them. The dissertation improved because I still had enough notes to rebuild. The present notebooks are gone more completely, so the evidentiary boundary must be correspondingly stricter.

Audit: a dissertation lesson used again

Finish what can be certified. Preserve the questions that exceed the present proof. Do not force an unfinished global result to authenticate a controlled local one. The current note follows the same rule John gave me in 2007.

10 The modern fiber view

10.1 Physical ranges over a base of choices

Collect the choices entering a constrained quantization into a base point

$$b = (O, \mu, t, \Lambda, \delta, n, \pi, \ldots) \in B, \tag{54}$$

where O may be a spacetime region, μ a scale, t time, Λ a regulator, δ the spectral window, n a gauge direction or auxiliary choice, and π a representation. At each b define

$$P_b = \chi_{[0,\delta_b^2]}(C_b), \qquad \mathcal{H}_b^{\text{phys}} = \operatorname{Ran} P_b. \tag{55}$$

If all C_b act on a common $\mathcal{H}$, depend continuously on b in norm resolvent sense, and a uniform gap separates the chosen spectral cluster, then $b \mapsto P_b$ is norm continuous. Locally, the ranges form a Hilbert subbundle of the trivial bundle $B \times \mathcal{H}$.

The canonical Grassmann connection is

$$\nabla = P\mathrm{d}, \tag{56}$$

with parallel-transport generator along a path

$$[\dot{P}, P]. \tag{57}$$

Its curvature is

$$F = P(\mathrm{d}P \wedge \mathrm{d}P)P. \tag{58}$$

One parameter is enough to prove motion but not curvature. Curvature requires at least two independent directions, for example scale and region, regulator and time, or two gauge parameters.

10.2 Projector ownership

The symbol P becomes dangerous when several physicalization procedures are present. A Klauder window projector,

$$P_\delta^K = \chi_{[0,\delta^2]}(C), \tag{59}$$

is not automatically the same object as a BRST harmonic projector,

$$P^{\text{BRST}} = \chi_{\{0\}}(\Delta_{\text{BRST}}), \tag{60}$$

or a limiting thin projector, a regulator-transported projector, a local vacuum projector, or a representation-specific support projection.

Each projector has an owner: the operator, representation, window, carrier, and comparison rule that define it. Two projectors may have isomorphic ranges while representing different observable algebras. They may coincide at one base point and separate under transport. They may agree before a singular limit and cease to live on the same carrier afterward.

It is therefore useful to write at least schematically

$$P = P[C, \delta, \pi, \Lambda, O, t, \dots] \tag{61}$$

rather than treating P as a timeless piece of punctuation. The expanded notation is cumbersome, but it prevents an equivalence at one layer from being silently promoted to equivalence at every layer.

A projector ownership checklist

Before comparing two projections, state: (i) their carrier Hilbert spaces; (ii) the self-adjoint operators or quadratic forms that generate them; (iii) the spectral windows; (iv) the representations of the observable algebra; (v) the regulator and limiting prescription; and (vi) the map by which their ranges are being identified.

10.3 When the carriers themselves vary

The simplest Grassmannian picture assumes that every P_b acts on one fixed Hilbert space. Field theory often violates that assumption. Different regions can carry different local algebras. Different infrared sectors can belong to inequivalent representations. Regulator removal may change the natural completion. A reduced reproducing kernel may define a physical space whose norm is not inherited from the original auxiliary carrier.

In such a regime, the comparison object need not be a unitary operator between subspaces of one $\mathcal{H}$. It may be a partial isometry, an isometric embedding, a correspondence, or a Hilbert W^*-module. The phrase "the projector moves" then has to be generalized: what moves is the physicalization package and the admissible comparison between its carriers.

This is also where the early reduced-kernel work returns. A family of projections can converge strongly to zero on the old Hilbert space while a renormalized kernel produces a nonzero new physical carrier. The failure of the old projector to survive is not the failure of every physical limit. It is evidence that the limiting room is not a subroom of the original building with the original norm.

10.4 Time dependence and effective evolution

If $P(t)$ varies, the naive compressed Hamiltonian $P(t)H(t)P(t)$ does not by itself keep a state inside the moving range. A natural effective generator is

$$H_{\text{eff}}(t) = P(t)H(t)P(t) + i[\dot{P}(t), P(t)], \tag{62}$$

under conventions where the Kato term transports the subspace. This is a Modern reading of the time-dependent issue, not a claim that equation (62) was written in the lost 2007 notes. It explains why a moving physical condition contributes to dynamics even when the ambient Hamiltonian is known.

10.5 Gap loss, rank change, and spectral flow

The bundle picture has a boundary. When the spectral gap closes, P_b need not be norm continuous. Rank can jump, discrete eigenvalues can meet a continuum, and a finite window can exchange states with its complement. The appropriate language is then spectral flow, continuous fields of Hilbert spaces, or convergence of quadratic forms rather than an ordinary vector bundle.

This boundary is physically interesting. Renormalization may drive the window and the operator together toward a singular limit. Region inclusion may change the carrier. Infrared constraints may prevent a common ambient Hilbert space. In such cases, comparison maps may be partial isometries or correspondences rather than unitaries.

The modern interpretation

The projection operator is not merely an answer at one point. A family $\{P_b\}$ records how physical state spaces sit inside, or correspond across, auxiliary carriers. Its derivative records motion, its curvature records comparison defect around infinitesimal loops, its holonomy records global failure of path-independent identification, and gap loss marks the failure of the smooth-bundle regime.

11 From fibers to a general comparison package

11.1 The physicalization package

For each base point b, write the data relevant to constrained physics as

$$\mathfrak{P}_b = \Big(\mathcal{H}_b, C_b, \delta_b, P_b, \mathcal{A}_b^{\mathrm{phys}}, \Omega_b\Big), \tag{63}$$

where $\mathcal{A}_b^{\mathrm{phys}}$ is the physical observable algebra and Ω_b a distinguished state when one exists. A comparison from b to b' must say more than "the Hilbert spaces are isomorphic." It should transport the physical range, observables, state, and dynamics to the degree claimed.

When all carriers live in one Hilbert space, a partial isometry may suffice. When representations differ, a Hilbert W^*-module or correspondence can be the right bridge. When only forms converge, Mosco-type convergence may replace operator-norm convergence. The most general categorical language is not developed here; what matters is that the comparison is itself physical data.

11.2 Motion is not curvature

A one-parameter family $P(s)$ can have $\dot{P}(s) \neq 0$ everywhere and still carry no curvature, because there is no two-dimensional infinitesimal loop on which a curvature two-form can be evaluated. Motion records that the physical range changes along a path. Curvature records path dependence when at least two independent directions are available.

For parameters (s, t),

$$F_{st} = P[\partial_s P, \partial_t P]P. \tag{64}$$

A nonzero F_{st} means that transporting first in the s direction and then in the t direction differs infinitesimally from transporting in the reverse order. Even when $F = 0$ locally, a non-simply-connected base can support nontrivial holonomy. Thus the hierarchy is

$$\text{motion} \not\Rightarrow \text{curvature}, \qquad \text{local flatness} \not\Rightarrow \text{global triviality}. \tag{65}$$

This hierarchy prevents the finite anomalous model from being asked to prove too much. It certifies motion along one deformation parameter. A curvature claim requires a second controlled parameter and a corresponding calculation.

11.3 Physical twisting

Suppose local physical ranges $P_i\mathcal{H}_i$ are defined over patches U_i of a base. On overlaps $U_i \cap U_j$, comparison maps

$$U_{ij} : P_j\mathcal{H}_j \longrightarrow P_i\mathcal{H}_i \tag{66}$$

may exist. On triple overlaps, one asks whether

$$U_{ij}U_{jk}U_{ki} = \mathbf{1} \tag{67}$$

on the physical range. Failure may be represented by a phase, a central operator, or a more general non-Abelian defect depending upon the setting.

The important methodological point is to keep the primary comparison object before collapsing it to one cohomology class. Different witnesses can detect different shadows of the same twisting: curvature, holonomy, spectral flow, determinant-line phase, or failure of observable intertwining. No single witness should be assumed complete without proof.

This is one reason the positive local range is valuable. It provides a place on which transition maps and their defects act. A global quotient can be correct while making this intermediate geometry difficult to see.

11.4 The windowed view

The base of a constrained theory may be stratified by spectral windows. On a region where a cluster remains isolated, P_b forms a smooth bundle. At a crossing, the rank or module type may change. In a singular limit, the old carrier may be replaced by a reduced-kernel space. One should therefore expect an atlas of valid comparison regimes rather than one globally rigid Hilbert room.

Within a window, ordinary differential geometry may apply. Between windows, comparison can require partial isometries, inclusions, correspondences, or limiting functors. The resulting theory is "windowed" not because physical reality is arbitrary, but because each mathematical representation has a declared regime of admissibility.

This modern view returns to Klauder's finite spectral band. The thickness δ was never merely a regulator to be erased without thought. It declared which part of the constrained spectrum was being treated as physical. Once several windows and carriers are compared, the declaration becomes part of the geometry.

Audit: strict comparison is earned

To replace a windowed family by one global physical Hilbert space, prove spectral persistence, compatible transport, flatness, trivial holonomy, observable and dynamical intertwining, and the absence of higher comparison defects. Abstract isomorphism of individual fibers does not earn the global replacement.

11.5 A strict global comparison test

In the smooth common-carrier regime, a proposed connection A on the physical bundle should satisfy

$$D_A P = 0, \qquad F_A = 0, \qquad \Omega_A = 0, \tag{68}$$

where $D_A P = 0$ expresses compatibility with the physical ranges, $F_A = 0$ flatness, and $\Omega_A = 0$ vanishing of any higher coherence defect in the chosen comparison scheme. These local conditions are not enough for a global identification. One must also require trivial holonomy and persistence of the relevant spectral cluster.

This is the "windowed view." There is no single once-and-for-all physical Hilbert space unless the comparison data prove that one can be assembled. The theory may instead provide a controlled family whose windows overlap, whose ranks persist on some strata, and whose transitions carry measurable defects.

Audit: one cohomology class is not a full comparison theory

A cohomology class can witness a local obstruction. It need not encode spectral persistence, positivity, region inclusions, states, dynamics, and holonomy simultaneously. The appropriate object is generally a witness family together with its coherence relations.

12 What the reconstruction establishes

The memoir has now reached the place where every major claim can be stated in one ledger. This section is intentionally more formal than the historical sections. A reader should be able to disagree with the modern interpretation without losing track of which mathematical statements have actually been proved.

12.1 The six-step mathematical spine

The note uses the following sequence.

1. Package the component constraints into a vector-valued operator $\mathbf{\Phi}_b$.
2. Form the positive squared constraint $C_b = \mathbf{\Phi}_b^\dagger G_b \mathbf{\Phi}_b$.
3. Select a low spectral range $P_b = \chi_{[0,\delta_b^2]}(C_b)$.
4. Test motion with the off-diagonal block $(1 - P_0)C_1 P_0$.
5. When a gap persists, transport the ranges with the Grassmann–Kato connection $P\mathrm{d}$.
6. When the gap or carrier fails, seek a reduced kernel or another admissible limiting carrier rather than assuming that the zero strong limit of the old projectors is the end of the theory.

The sequence connects the 2004 kernel question, the 2005 anomalous model, the 2006–2007 loop question, and the present fiber interpretation. It is not a claim that all four stages were formulated in this language at the time.

12.2 Certified claims

The following claims are certified within the scope stated in the note.

C1. A closed positive squared-constraint form has zero kernel equal to the common kernel of the vector constraint, under the positivity hypotheses of (8).

C2. A pointwise limit of positively rescaled positive kernels is positive and therefore defines a reproducing-kernel Hilbert space after null-vector quotient and completion.

C3. Singular rescaling does not automatically preserve idempotence with respect to the original coherent-state measure.

C4. The finite anomalous model has a constant-rank physical line with nonzero Fubini–Study and Hilbert–Schmidt speeds. Its projector genuinely moves.

C5. For an isolated spectral cluster, first-order range motion is equivalent to a nonzero off-diagonal block $(1 - P_0)C_1 P_0$.

C6. Under closed-range Hodge hypotheses, BRST cohomology is represented by the harmonic subspace; under an additional gap hypothesis, that space equals a finite-window spectral range of the BRST Laplacian.

C7. Full equivalence between BRST and a separately constructed Klauder physicalization requires an intertwiner of observables, dynamics, and transport in addition to a vector-space isomorphism.

12.3 The claims that remain historical

The following statements are used as remembered chronology rather than as proofs of mathematical priority:

- Wayne Bomstad and Sergei Shabanov were pursuing projection-operator QED with interacting scalar and spinor matter.
- Wayne's loop calculation was encountering difficulty before I left Florida in 2006.
- During quarterly visits we discussed the calculation and possible techniques while preserving Wayne's ownership of his route.
- I pursued an adjacent Yang–Mills calculation using the treatment of algebraic noncovariant gauges in the Bassetto–Nardelli–Soldati text.
- The completed notebook calculation appeared to require a correction to the squared-constraint physical projection.
- I showed John an early BRST comparison and time-dependent results in May 2007, while judging the Yang–Mills work not ready for the dissertation.

These memories explain why the questions were asked. They do not substitute for the missing operator calculation.

12.4 What is not claimed

The boundaries are equally important.

This note does not recover the exact one-loop Yang–Mills master operator. It does not determine which precise member of the algebraic noncovariant gauge family was used. It does not prove that the renormalized witness $W_{\Lambda,1}$ was nonzero. It does not construct interacting continuum Yang–Mills theory, establish confinement, prove a mass gap, or settle the global infrared problem.

It does not claim that BRST is generally defective. It identifies the domain on which a Hodge representative space yields a positive spectral realization and lists the additional data needed for comparison with another physicalization.

It does not claim that every moving projector carries curvature. A one-parameter family can move while remaining flat. Curvature requires at least two controlled directions and a comparison around infinitesimal loops.

It does not claim that every vanishing family of projections has a nontrivial reduced-kernel limit. The existence and normalization of such a limit have to be proved.

Finally, it does not attribute the modern GAT framework to Klauder. The modern framework is being used to read an operator architecture whose physicalizing step Klauder deliberately made visible.

12.5 The sharper questions

The reconstruction replaces the old question, "Why did my calculation not close?" with a list of testable questions:

1. Is $(1 - P_{0,\Lambda})C_{1,\Lambda}P_{0,\Lambda}$ nonzero?
2. If the first-order witness vanishes, what is the first perturbative order at which the physical range moves?
3. Does the range rotate while its lowest eigenvalue remains zero, or does the low eigenvalue also lift?
4. Must the physical window $\delta_{g,\Lambda}$ run with coupling and scale?
5. Is the motion invariant under the allowed counterterms and regulator changes?
6. Do different algebraic noncovariant gauges produce equivalent physical transport?
7. Does the family of physical projectors converge as the regulator is removed? If not, does a rescaled reproducing kernel converge?
8. Is the BRST harmonic family intertwined with the Klauder family at the levels of observables, dynamics, and transport?
9. Does the two-parameter family $P_{g,\Lambda}$ carry nonzero mixed curvature?
10. Can the local physical ranges be assembled globally without holonomy or higher coherence defect?

These questions can be wrong, answered negatively, or rendered ill-posed by a more careful reconstruction. Their value is that each has a mathematical witness. The lost calculation has been converted from a remembered outcome into a falsifiable program.

The strongest defensible historical statement

Between 2004 and 2007, the author studied reduced reproducing kernels, singular projection limits, anomalous finite constraints, time-dependent physicalization, BRST comparison, and perturbative corrections to squared-constraint projections. The finite model certifies nontrivial projector motion. A related Yang–Mills loop calculation was attempted but remains unrecovered. Modern spectral and fiber geometry show that the remembered motion is mathematically natural and may coexist with closure and constant physical rank.

13 Why Klauder is the G.O.A.T. on this mountain

13.1 What the method got right

The joke works because the mountain is real. Klauder's projection-operator program made a series of decisions that age remarkably well:

1. **Quantize before reducing.** The quantum constraint operator, not only the classical zero set, determines the physical range.
2. **Use positivity.** The sum of squared constraints supplies a self-adjoint spectral problem and a positive physical inner product.
3. **Keep a nonzero window.** First-class, second-class, anomalous, irregular, and continuous-spectrum problems can be discussed in one operator language.
4. **Do not require a gauge slice.** Gauge fixing and ghosts may be useful, but they are not prerequisites for defining the spectral range.
5. **Expose the physicalizing map.** Reduced kernels and projections show how the physical carrier is produced and therefore how it can change.

The last point is the one this note needed. A formal quotient can conceal the motion of representatives. A projection displays it. Once displayed, the motion can be differentiated, compared, and assigned curvature.

13.2 Praise with a boundary

Calling Klauder the G.O.A.T. here is not a claim that the projection method has already solved interacting Yang–Mills theory, the global infrared problem, renormalization of every master constraint, local quantum field theory, or the comparison of all representations. It is a judgment about architecture. He put the positive physical range in the place where later geometry can find it.

The finite anomalous example also shows the scientific virtue of a well-chosen small problem. It certified the possibility that the physical range changes under quantum deformation. The highly irregular examples certified that different physical sectors may require different limiting scales. The field examples certified that the method was not confined to finitely many constraints. Together they created a path that the modern fiber view can now extend without rewriting the history.

13.3 He let the small problem be serious

One quality of John's work was his refusal to confuse a small model with a small question. A finite example could be worth doing if it isolated a structural difficulty that a field theory would bury beneath technique. The anomalous model did exactly that. It separated classical equivalence from quantum physicalization in a setting where every relevant vector could be written down.

This mattered to a young physicist. Grand problems can make every tractable calculation look like an evasion. John showed another possibility: choose a problem small enough that the obstruction cannot hide, then let the formal mathematics tell you what the obstruction actually is.

Wayne's excitement that "formal mathematics enters the fray" belongs to this atmosphere. The mathematics was not being added to decorate a physical argument. It entered because positivity, kernel reduction, ordering, and spectral selection had become the physics of the problem.

That lesson remained with me even when I underestimated these particular papers. Much of my later work follows the same pattern: construct a finite or regulated carrier on which the decisive obstruction is visible, certify it there, and then ask what structure is required to transport it toward the larger theory.

13.4 He did not require premature certainty

Shabanov's question — "What does this say about quantum field theory?" — was exactly the right pressure. My answer, "Nothing yet," was not a dismissal of the finite result. It recognized the distance between a certified mechanism and a field-theoretic theorem.

John's style made room for that distance. He could be encouraged by a time-dependent result or a BRST comparison without pretending that the Yang–Mills calculation was ready. He could say that a larger result might be real and still advise that it wait outside the dissertation.

That combination is rare and valuable. Skepticism without encouragement can stop a line of inquiry before it matures. Encouragement without boundaries can turn an intuition into a claim too early. John's response was closer to: there may be something here; learn what you can; finish what you can defend; keep moving.

The present note tries to preserve the same balance. The modern fiber view is allowed to show why the old puzzle mattered. The lost Yang–Mills result is not allowed to become certified merely because the modern view makes it plausible.

13.5 The Billy G.O.A.T. and the window

The joke contains a serious mathematical judgment. Klauder did not insist that every constraint problem be forced through the same classical reduction or exact-zero prescription. He retained a window:

$$P_\delta = \chi_{[0,\delta^2]}(C). \tag{69}$$

That window can accommodate exact first-class constraints, partially second-class systems, anomalous finite models, continuous spectrum, and irregular sectors whose limits require different magnifications.

The window also prevents the physicalizing operation from disappearing behind a quotient. It tells us which spectral data are retained and on which carrier. Once written as an operator, it can be differentiated. Once differentiated, its motion can be measured. Once a family of such operators is considered, transport, curvature, holonomy, and singular limits become available questions.

Klauder did not write the modern GAT treatment contained here. The praise is more precise: he built an architecture in which the modern questions have an object on which to act. He put the door to the physical Hilbert space on its hinges.

13.6 What the mountain view changes personally

For a long time, I filed the early papers under work completed before the important work began. They were explicit, finite, and connected to a dissertation period whose larger ambitions were left unfinished. From that angle, they looked preliminary.

The mountain view does not turn them into a solution of Yang–Mills theory. It changes their position in the intellectual history. Reduced kernels asked how a physical carrier survives a singular limit. Irregular constraints asked whether different sectors require different limiting scales. The anomalous model showed that quantum physicalization can move even when the classical zero surface is unchanged. The time-dependent work asked how a changing physical condition enters evolution. The BRST comparison asked when a cohomology and a positive spectral range truly agree.

Those are seeds. A seed is not a mature tree, and saying so is not sentiment. It is a statement about the continuity of the questions. The later framework did not appear without a past. Some of the habits it requires were learned in those finite calculations at Florida.

That recognition is therapeutic because it corrects a false institutional story as well. Serious physics was not waiting somewhere else for me to arrive. It was happening in Gainesville, in John's office, in conversations with Wayne and Shabanov, in a

Christmas book opened for an adjacent problem, and in a dissertation rewritten after half of it vanished.

Mountain memory: The therapy of a correct retrospective

It touches me that the earliest work with Klauder belongs in this story without being forced into it. What once looked somewhat trivial now shows the seeds it planted. The lesson is not that every early calculation was secretly a grand theory. It is that I was doing physics, and the questions were already real. One did not have to be at an Ivy or MIT for that to be true. Florida was a very good school.

13.7 The view from the mountain

The old question was why the projection operator would need quantum corrections. The modern answer is that the operator being projected has changed, and the physical range is a spectral object, not a painted boundary. Sometimes the correction only recalibrates the window. Sometimes it rotates the range. Sometimes the gap closes and the very notion of a fixed-rank bundle fails. Each case says something different, and the projection method keeps them separate.

The BRST comparison leads to the same moral. When a Hodge theorem, a gap, and an intertwiner exist, cohomology and projection agree. When they do not, the failure should be localized: nonclosed range, continuous spectrum, second-class structure, or incompatible gluing. One global slogan should not erase these local diagnoses.

The unfinished Yang–Mills calculation therefore has a sharper status than either success or failure. It supplied the right physical discomfort before the right mathematical question had been isolated. The 2005 model proves that the phenomenon can occur. The one-loop reconstruction now has a target: calculate $W_{\Lambda,1}$, control the gap and counterterms, and test whether the resulting family carries nontrivial motion or curvature.

Closing claim

The physical state space is not always a single global room waiting to be entered. It can be a family of positive spectral ranges, stable on some windows, singular on others, and connected by transport that has physical content. Klauder taught us to build the door as an operator. Once the door is visible, we can finally see that it moves.

Every mountain needs a goat. This one had John Klauder.

A Surviving chronology and evidence ledger

The chronology below is intentionally conservative. Dates attached to publications are recorded; autobiographical transitions are remembered.

Date	Status	Event
2004–2005	Recorded	Reduced-kernel and anomalous-constraint work reaches publication with Klauder. The finite anomaly model contains a moving rank-one physical range.
2005	Remembered	Wayne reacts with excitement; Shabanov asks what the result says about quantum field theory.
2005–2006	Recorded	Highly irregular constraints and componentwise reductions are developed. Bomstad and Klauder complete the linearized-gravity projection analysis.
Christmas 2005	Remembered	Receipt of Bassetto, Nardelli, and Soldati's book motivates work in algebraic noncovariant gauges.
2006	Remembered	Wayne and Shabanov pursue interacting QED with scalar and spinor matter; loop calculations become difficult. The Yang–Mills adjacent problem begins.
2006–2007	Remembered	A one-loop squared-constraint Yang–Mills calculation is completed in a noncovariant gauge; its moving projection is not understood.
May 2007	Remembered	A BRST-equivalence argument and first time-dependent results are shown to Klauder. The Yang–Mills work is judged not ready for the dissertation.
July 2007	Remembered	A chapter-by-chapter dissertation outline is provisionally approved; the first three chapters are requested.
2007	Remembered	Roughly half the dissertation is lost and reconstructed from notes; the rewrite improves the result.
Present	Modern reading	Projector perturbation, Kato transport, curvature, and comparison packages supply the language for the old phenomenon.

B Constraint compression as a positive-form theorem

This appendix states carefully what it means to replace many constraints by one squared-constraint operator. The point is elementary in finite dimensions, but the field-theoretic use requires form domains and a distinction between exact and windowed physicalization.

B.1 The vector constraint

Let $\mathcal{H}$ be a Hilbert space and let $\mathcal{K}_{\mathrm{con}}$ be a Hilbert space of constraint labels. A finite-dimensional compact gauge algebra gives $\mathcal{K}_{\mathrm{con}} \cong \mathfrak{g}^*$; a regulated field theory includes spatial smearing labels as well. Let

$$\mathbf{\Phi} : D(\mathbf{\Phi}) \subset \mathcal{H} \longrightarrow \mathcal{H}\widehat{\otimes}\mathcal{K}_{\mathrm{con}} \tag{70}$$

be closed and densely defined. Its components relative to an orthonormal label basis are the operators Φ_a.

Let G be a bounded positive operator on $\mathcal{K}_{\mathrm{con}}$ satisfying

$$G \geq g_0 \mathbf{1}, \qquad g_0 > 0, \tag{71}$$

on the admitted label space. Define

$$q_C[\psi] = \langle \mathbf{\Phi}\psi, (\mathbf{1} \otimes G)\mathbf{\Phi}\psi \rangle, \qquad D(q_C) = D(\mathbf{\Phi}). \tag{72}$$

Because $\mathbf{\Phi}$ is closed and G is positive and bounded above and below, q_C is a densely defined closed nonnegative form. The representation theorem for closed forms gives a unique nonnegative self-adjoint operator C such that

$$q_C[\psi] = \|C^{1/2}\psi\|^2. \tag{73}$$

Theorem B.1 (Exact constraint compression)**.** *Under the preceding assumptions,*

$$\operatorname{Ker} C = \operatorname{Ker} C^{1/2} = \operatorname{Ker} \mathbf{\Phi}. \tag{74}$$

If the components Φ_a are defined by the label basis, then

$$\operatorname{Ker} C = \bigcap_a \operatorname{Ker} \Phi_a. \tag{75}$$

Proof. If $\mathbf{\Phi}\psi = 0$, then $q_C[\psi] = 0$, hence $C^{1/2}\psi = 0$. Conversely, if $C^{1/2}\psi = 0$, then $q_C[\psi] = 0$. The lower bound (71) implies

$$0 = q_C[\psi] \geq g_0 \|\mathbf{\Phi}\psi\|^2, \tag{76}$$

so $\mathbf{\Phi}\psi = 0$. For a vector-valued operator, vanishing of the vector is equivalent to vanishing of every component. □

This theorem is the rigorous content of "looking in one direction." The single direction is the nonnegative spectrum of C. The theorem does not identify the component operators with one another or remove the non-Abelian relations among them.

B.2 What changes when G is only semidefinite?

If G has a kernel, then q_C does not see every constraint-label direction. One has only

$$\operatorname{Ker} \mathbf{\Phi} \subseteq \operatorname{Ker} C, \tag{77}$$

and the inclusion may be strict. A constraint vector can lie entirely in the null space of G without vanishing. Thus the positivity used in the master constraint is not a formal preference; it decides whether any constraint directions have been discarded.

Similarly, an indefinite contraction does not define a norm of constraint violation. Cancellations between positive and negative directions can make the quadratic expression vanish while individual constraints remain nonzero. The spectral-window interpretation depends upon nonnegativity.

B.3 Exact kernel versus nonzero window

The exact theorem does not require a spectral gap. It identifies the kernel, whether or not zero is isolated. Klauder's physicalization uses more:

$$P_\delta = \chi_{[0,\delta^2]}(C). \tag{78}$$

For $\psi \in \operatorname{Ran} P_\delta$,

$$\langle \psi, C\psi \rangle \leq \delta^2 \|\psi\|^2, \tag{79}$$

so

$$\|G^{1/2}\mathbf{\Phi}\psi\| \leq \delta \|\psi\|. \tag{80}$$

The physical window therefore controls the total constraint norm. It does not generally imply the pointwise component estimate $\|\Phi_a\psi\| \leq \delta\|\psi\|$ with a uniform constant independent of the label geometry; such a conclusion requires bounds relating the components to the G-norm.

The nonzero window is indispensable for second-class systems and for problems in which zero is not a normalizable eigenvalue. It also turns the window thickness into physicalization data. Two investigators using the same C but different δ need not be discussing the same physical range.

B.4 Compact-group representation decomposition

Let a compact group G_c act unitarily on $\mathcal{H}$. The Peter–Weyl-type decomposition takes the schematic form

$$\mathcal{H} \cong \widehat{\bigoplus}_{R\in\widehat{G}_c} \mathcal{M}_R \otimes V_R. \tag{81}$$

For an invariant inner product on the Lie algebra, the quadratic Casimir is central in the enveloping algebra. Schur's lemma gives

$$C_2|_{\mathcal{M}_R\otimes V_R} = c_2(R)\mathbf{1}_{\mathcal{M}_R\otimes V_R}. \tag{82}$$

For compact semisimple groups, $c_2(R) = 0$ only for the trivial representation. Hence

$$\operatorname{Ker} C_2 \cong \mathcal{M}_{\mathrm{triv}}. \tag{83}$$

The master operator does not tell us merely how many singlets occur. The spectral multiplicities record which nontrivial representation sectors lie at each Casimir value. Degenerate Casimir eigenvalues can contain inequivalent representations, so C_2 alone need not reconstruct the full representation ring. This is another useful boundary: compression preserves enough information to select singlets and organize sectors, but it need not be a complete invariant of representation content.

B.5 Local Gauss law

For a regulated Yang–Mills theory, let $G_a(f)$ denote a Gauss constraint smeared with a test function f. A positive kernel on color and smearing labels gives schematically

$$C_\Lambda = \sum_{a,b} \int \mathrm{d}x\, \mathrm{d}y\; G_a(x)^\dagger K_\Lambda^{ab}(x,y) G_b(y). \tag{84}$$

At finite regulator, one seeks a closed form analogous to (72). The continuum question is not answered by writing (84); one must control the limit of the forms, their domains, and their low spectral ranges.

The field-theoretic advantage nevertheless remains visible. Instead of a separate spectral analysis for every a and every f, the regulated theory produces one nonnegative operator C_Λ. The cost is that regulator dependence is concentrated into that operator and its physical window.

B.6 Perturbing the vector constraint

Let

$$\mathbf{\Phi}(g) = \mathbf{\Phi}_0 + g\mathbf{\Phi}_1 + O(g^2) \tag{85}$$

on a common perturbative form domain. The first correction to the squared operator is given by (17). If P_0 projects onto $\operatorname{Ker} \mathbf{\Phi}_0$, then

$$P_0 C_1 P_0 = 0. \tag{86}$$

Indeed, both terms contain one factor of $\mathbf{\Phi}_0$ adjacent to a physical vector or covector. Equation (86) explains why a first-order eigenvalue shift can vanish automatically for an exact squared constraint.

But the off-diagonal block is

$$(1 - P_0) C_1 P_0 = (1 - P_0) \mathbf{\Phi}_0^\dagger G \mathbf{\Phi}_1 P_0, \tag{87}$$

which need not vanish. Thus the representation sector can tilt before its spectral value lifts. Squaring compresses the directions; it does not freeze their embedded common kernel.

Compression audit

A valid use of the squared constraint should state: the vector constraint and its form domain; the positive label metric; the associated self-adjoint operator; the exact or windowed physical condition; the representation and regulator; and which of these data vary in the comparison being made.

C Reduced-kernel theorem in a reusable form

Theorem C.1 (Positive reduced kernel)**.** *Let X be a set and K_ϵ a positive kernel on X for every $\epsilon > 0$. Let $s_\epsilon > 0$. If*

$$K(x, y) = \lim_{\epsilon \downarrow 0} s_\epsilon K_\epsilon(x, y) \tag{88}$$

exists for all $x, y \in X$, then K is positive. If K is not identically zero, it defines a nontrivial RKHS.

Proof. For each finite $x_1, \ldots, x_n$ and $c \in \mathbb{C}^n$, positivity gives

$$s_\epsilon \sum_{i,j=1}^{n} \bar{c}_i c_j K_\epsilon(x_i, x_j) \geq 0. \tag{89}$$

The finite sum commutes with the pointwise limit, so the corresponding quadratic form for K is nonnegative. Quotient the span of kernel sections by the null space and complete. □

Remark C.2. No fixed measure is needed for the conclusion. If each K_ϵ is the kernel of an orthogonal projection in one $L^2(X, \mu)$, singular rescaling can destroy idempotence relative to μ even while preserving positivity. The limiting RKHS, not the old integral identity, is the invariant object.

Proposition C.3 (Orthogonal component reduction)**.** *Suppose*

$$K_\epsilon = \sum_{a=1}^{r} K_\epsilon^{(a)} + R_\epsilon \tag{90}$$

where the sectors are mutually orthogonal in the pre-limit Hilbert spaces and $s_a(\epsilon) K_\epsilon^{(a)} \to K^{(a)}$ pointwise. If the rescaled remainders vanish in all sector pairings, the reduced physical space is the Hilbert direct sum of the RKHSs of the nonzero $K^{(a)}$.

The proposition is less a trick than a warning. A dominant asymptotic sector can hide other physical components. The correct limiting procedure must be sensitive to the spectral and kernel decomposition already present.

C.1 Why the limit is called a fixed physical point

For every $\epsilon > 0$, the kernel K_ϵ describes a physical range at finite window or finite regulator. The rescaling

$$K_\epsilon \longmapsto s_\epsilon K_\epsilon \tag{91}$$

changes the normalization of the physical inner product. If the rescaled kernel converges to K, then K is unchanged by continuing the limiting prescription: it is a fixed output of the reduction rather than another vanishing representative on the original scale.

The word "fixed" should not be confused with a fixed vector inside the auxiliary Hilbert space. The limiting RKHS may have a new norm and may not embed continuously into the original carrier. What is fixed is the positive kernel data after the required magnification.

This distinction anticipates the later fiber view. A singular limit need not produce a subspace of the original Hilbert space. It may produce a new fiber whose comparison with the pre-limit fibers is mediated by rescaled kernel sections.

C.2 Why old-measure idempotence can fail

Before rescaling, suppose

$$\int_X K_\epsilon(x,z)K_\epsilon(z,y)\,\mathrm{d}\mu(z) = K_\epsilon(x,y). \tag{92}$$

After multiplying by s_ϵ,

$$\int_X [s_\epsilon K_\epsilon(x,z)][s_\epsilon K_\epsilon(z,y)]\,\mathrm{d}\mu(z) = s_\epsilon[s_\epsilon K_\epsilon(x,y)]. \tag{93}$$

If s_ϵ diverges, the right side has an extra divergent factor. Therefore the rescaled limit cannot generally satisfy the same idempotence identity with the same measure.

This does not damage positivity. A positive kernel defines its native Hilbert space through finite linear combinations of kernel sections,

$$f = \sum_i c_i K(\,cdot\,, x_i), \tag{94}$$

with inner product

$$\langle f, g\rangle_K = \sum_{i,j} \overline{c_i} d_j K(x_i, y_j). \tag{95}$$

The reproducing property then follows from the RKHS construction itself. What has been lost is only the representation of that property as the old integral projector on $L^2(X,\mu)$.

C.3 Several asymptotic scales

Suppose two orthogonal sectors behave as

$$K_\epsilon^{(1)} \sim \epsilon K^{(1)}, \qquad K_\epsilon^{(2)} \sim \epsilon^3 K^{(2)}. \tag{96}$$

Multiplication by ϵ^{-1} recovers the first sector and sends the second to zero. Multiplication by ϵ^{-3} recovers the second but makes the first diverge. No single scalar normalization recovers both as a finite direct sum.

The componentwise prescription uses the spectral or kernel decomposition before taking the limit:

$$\epsilon^{-1} K_\epsilon^{(1)} o K^{(1)}, \qquad \epsilon^{-3} K_\epsilon^{(2)} o K^{(2)}. \tag{97}$$

The physical limit is then

$$\mathcal{H}_K = \mathcal{H}_{K^{(1)}} \oplus \mathcal{H}_{K^{(2)}}. \tag{98}$$

The different rates are not nuisance parameters. They record how separate physical sectors approach the singular surface.

C.4 What the theorem does not decide

The positivity theorem does not select the normalization s_ϵ. That choice must come from the asymptotics, a physical normalization condition, or comparison with a known sector. Different normalizations can produce unitarily equivalent limits, inequivalent limits, or a trivial zero kernel.

Nor does pointwise convergence automatically provide convergence of observables, dynamics, or measures. A complete limiting theory needs compatible convergence of whatever additional structures are claimed. The kernel theorem certifies the physical carrier; it does not silently transport the entire theory.

Audit: reduced-kernel claim boundary

Proved: positivity of the rescaled pointwise limit and existence of its RKHS. Not automatic: uniqueness of scale, old-measure idempotence, embedding into the auxiliary carrier, convergence of observables, or equivalence of distinct componentwise reductions.

D The anomalous finite model as a Grassmannian calculation

Let

$$e_1 = |200\rangle, \qquad e_2 = |020\rangle, \qquad e_3 = |002\rangle, \tag{99}$$

and $u = (e_1 + e_2 + e_3)/\sqrt{3}$. From (32),

$$v(\beta) = \frac{(1-2\beta)e_1 + (1-\beta)e_2 + e_3}{\sqrt{(1-2\beta)^2 + (1-\beta)^2 + 1}} + O(\beta^2). \tag{100}$$

Since the denominator is $\sqrt{3}(1-\beta) + O(\beta^2)$,

$$v(\beta) = u + \beta w + O(\beta^2), \qquad w = \frac{-e_1 + e_3}{\sqrt{3}}, \qquad \langle u, w\rangle = 0. \tag{101}$$

For $P(\beta) = |v(\beta)\rangle\langle v(\beta)|$,

$$\dot{P}(0) = |w\rangle\langle u| + |u\rangle\langle w|. \tag{102}$$

Therefore

$$\mathrm{Tr}\Big(\dot{P}(0)^2\Big) = 2\langle w, w\rangle = \frac{4}{3}. \tag{103}$$

The Fubini–Study line element is

$$\mathrm{d}s^2 = \Big(\langle \mathrm{d}v, \mathrm{d}v\rangle - |\langle v, \mathrm{d}v\rangle|^2\Big) = \frac{2}{3}\,\mathrm{d}\beta^2 + O(\beta). \tag{104}$$

The calculation is gauge invariant under $v(\beta) \mapsto e^{i\theta(\beta)}v(\beta)$ and depends only on the line. It is therefore an intrinsic certificate that the physical projector moves.

D.1 Principal angle and projector distance

For normalized vectors u and $v(\beta)$, the principal angle between their lines is

$$\theta(\beta) = \arccos|\langle u, v(\beta)\rangle|. \tag{105}$$

Since $v(\beta) = u + \beta w + O(\beta^2)$ with $w \perp u$, normalization gives

$$|\langle u, v(\beta)\rangle|^2 = 1 - \beta^2\|w\|^2 + O(\beta^3). \tag{106}$$

Therefore

$$\theta(\beta) = |\beta|\sqrt{\frac{2}{3}} + O(\beta^2). \tag{107}$$

For rank-one projections,

$$\|P(\beta) - P(0)\|_{\mathrm{HS}}^2 = 2\left(1 - |\langle u, v(\beta)\rangle|^2\right) = \frac{4}{3}\beta^2 + O(\beta^3). \tag{108}$$

Thus the derivative calculation, Fubini–Study metric, principal angle, and Hilbert–Schmidt projector distance all give the same nonzero motion witness.

D.2 Constant rank is part of the point

The example is not a rank-jump phenomenon. Both physical projectors have rank one for the small deformation under discussion. The physical degree of freedom count is unchanged. What changes is the embedding of that one dimensional physical space in the auxiliary carrier.

This is why an argument based only on counting states would miss the result. Two one-dimensional Hilbert spaces are abstractly unitarily isomorphic. The constraint problem contains more information: which line is selected, how observables compress to it, and how that line is compared along the deformation.

D.3 The lifted eigenvalue and retained window

The low eigenvalue behaves as

$$\lambda_{\min}(\beta) \sim \frac{1}{2}\hbar^2\beta^6, \tag{109}$$

while the chosen window has scale

$$\delta^2 = \beta\hbar^2. \tag{110}$$

For sufficiently small positive β, the lifted state remains well inside the physical band. The projector is therefore not the exact zero projection of the modified master operator; it is the low-band physical projector appropriate to the partially second-class system.

The model consequently certifies two different departures from a naive exact Dirac picture. The physical line rotates, and the admissible physical state need not lie at exact spectral zero. Klauder's window handles both without pretending they are the same effect.

Certified finite conclusion

The model proves constant-rank Grassmannian motion of a window-selected physical sector under a quantum constraint deformation. It does not prove the analogous statement for Yang–Mills; it proves that such a conclusion is mathematically coherent and supplies the geometry with which to state it.

E Spectral-projector perturbation theory

Let $C(s) = C_0 + sC_1 + O(s^2)$ be a self-adjoint analytic family in a regime where a compact spectral cluster σ_0 is separated from the rest by a contour Γ. Then

$$P(s) = \frac{1}{2\pi i}\oint_\Gamma (z - C(s))^{-1}\,\mathrm{d}z \tag{111}$$

is analytic. Differentiation gives (42). Decomposing relative to P_0 yields

$$P_0P_1P_0 = (1-P_0)P_1(1-P_0) = 0. \tag{112}$$

Thus P_1 is entirely off-diagonal, as required by differentiating $P^2 = P$.

If the low cluster is an eigenvalue λ_i and the complement has eigenvectors $|a\rangle$, differentiating $[C, P] = 0$ gives

$$(\lambda_a - \lambda_i)\langle a|P_1|i\rangle = -\langle a|C_1|i\rangle. \tag{113}$$

The reduced resolvent form is

$$P_1 = -SC_1P_0 - P_0C_1S, \qquad S = (1-P_0)(C_0-\lambda_i)^{-1}(1-P_0), \tag{114}$$

for a simple isolated eigenvalue. The diagonal part $P_0C_1P_0$ shifts the eigenvalue inside the cluster. The off-diagonal part rotates the cluster.

One-loop audit protocol

1. Specify the regulated self-adjoint operator $C_{\Lambda,R}$ and its form domain.
2. Specify the renormalized window $\delta_R(\Lambda)$.
3. Prove a separating gap or replace analytic perturbation by an appropriate continuous-spectrum method.
4. Compute $(1-P_0)C_1P_0$ after all allowed counterterms.
5. Test regulator removal and gauge-parameter dependence of the resulting transport.
6. Only then claim range motion, curvature, or a physical anomaly.

E.1 The Riesz derivative without an eigenbasis

The contour formula is preferable to an eigenvector expansion when the low cluster is degenerate. Let

$$R_s(z) = (z - C(s))^{-1}. \tag{115}$$

Differentiating the inverse gives

$$\dot{R}_0(z) = R_0(z)C_1R_0(z), \tag{116}$$

and therefore

$$P_1 = \frac{1}{2\pi i}\oint_\Gamma R_0(z)C_1R_0(z)\,\mathrm{d}z. \tag{117}$$

No choice of basis inside the cluster is involved. Internal degeneracy may make the eigenvectors chosen by ordinary perturbation theory change abruptly, while the spectral projection remains smooth. The Riesz projection follows the subspace rather than a preferred frame within it.

Differentiating $P(s)^2 = P(s)$ yields

$$P_0P_1P_0 = 0, \qquad (1 - P_0)P_1(1 - P_0) = 0. \tag{118}$$

Every tangent vector to the Grassmannian is off-diagonal. The diagonal block $P_0C_1P_0$ can split energies within the selected cluster, but it does not by itself move the range of the cluster.

For a zero cluster with $C_0 \geq \gamma(1 - P_0)$, define the reduced inverse

$$S = (1 - P_0)C_0^{-1}(1 - P_0). \tag{119}$$

Then

$$P_1 = -SC_1P_0 - P_0C_1S. \tag{120}$$

Because S is injective on the complementary spectral space,

$$P_1 = 0 \iff (1 - P_0)C_1P_0 = 0. \tag{121}$$

E.2 Four phenomena that should not share one name

A statement that "the constraint receives a quantum correction" is too coarse. At least four different spectral events are possible.

Eigenvalue displacement. The low eigenvalues may move while the spectral subspace remains fixed. If $C(g)$ is diagonal in a g-independent decomposition, then $P(g)$ can remain constant even though every selected eigenvalue changes. This can require a running window without producing range motion.

Constant-rank range motion. The rank remains fixed and the cluster remains isolated, but $P(g) \neq P(0)$. This is ordinary smooth Grassmannian motion. It is the phenomenon certified by the finite anomalous model.

Rank change. An eigenvalue crosses a window boundary or a zero mode is created or removed. Then rank $P(g)$ changes. A smooth finite-rank vector bundle cannot cross such a point without stratification.

Loss of spectral isolation. The gap closes, a discrete cluster meets continuous spectrum, or zero becomes an accumulation point. The contour Γ can no longer be held fixed. The analytic Riesz calculation has reached its boundary of validity.

These events can occur together. They should nevertheless be diagnosed separately because they require different mathematical tools and carry different physical meanings.

E.3 First nonvanishing order

The first-order witness may vanish while the projector moves later. Write

$$C(g) = C_0 + gC_1 + g^2C_2 + O(g^3), \qquad P(g) = P_0 + gP_1 + g^2P_2 + O(g^3). \tag{122}$$

If $P_1 = 0$, the second-order term depends both on C_2 and on repeated excursions generated by C_1 through the complementary space. Schematically, the effective cluster operator contains

$$P_0C_2P_0 - P_0C_1SC_1P_0. \tag{123}$$

The corresponding range derivative contains off-diagonal terms built from C_2 and quadratic combinations of C_1.

Therefore a vanishing one-loop witness proves only stationarity at that order. It does not prove that the physical range is fixed to all orders. This warning is especially relevant for squared constraints, whose positivity can force some diagonal spectral shifts to begin at even order.

E.4 Window motion is separate data

Let

$$P(g) = \chi_{[0,\delta(g)^2]}(C(g)). \tag{124}$$

Even if the eigenspaces of $C(g)$ are constant, changing $\delta(g)$ can admit or exclude states. Conversely, an isolated cluster can rotate while every member remains comfortably within a fixed window. The derivative of the indicator function is singular at a crossing, so one should not treat δ as an ordinary smooth parameter when spectral values touch the boundary.

The clean procedure is to select a cluster by a contour as long as a gap persists, analyze its range motion, and only then ask whether the physical window continues to select that cluster. This separates the geometry of the cluster from the policy by which it is admitted as physical.

E.5 Counterterms and changes of identification

A correction to C can sometimes be removed by an admissible unitary change of carrier-space coordinates. If

$$C(g) = U(g)C_0U(g)^\dagger \tag{125}$$

for a controlled unitary family, then

$$P(g) = U(g)P_0U(g)^\dagger. \tag{126}$$

The projector moves in the fixed trivialization, but the family is unitarily equivalent to a constant one. Whether that motion is physically trivial depends upon whether $U(g)$ also intertwines the observable algebra, dynamics, states, locality, and the chosen comparison structure.

Similarly, adding a counterterm proportional to C_0 can rescale spectral values without changing eigenspaces. An off-diagonal counterterm may rotate the range. The phrase "removed by renormalization" must specify which counterterms are admissible and which physical structures they preserve.

E.6 When the gap closes

If C_Λ has a finite-regulator gap that closes as $\Lambda \to \infty$, the family P_Λ may fail to converge in norm. Several outcomes are possible:

$$P_\Lambda \to P \quad \text{strongly}, \qquad P_\Lambda \to 0 \quad \text{strongly}, \tag{127}$$

or no strong limit may exist. Weak convergence can retain matrix elements while losing norm control. A reduced kernel can survive even when the strong operator limit is zero.

The appropriate tools may include strong or norm resolvent convergence, Mosco convergence of quadratic forms, spectral flow, or convergence of reproducing kernels. The particular tool depends upon which structure is actually controlled. One should not cite analytic perturbation theory across a gap-closing limit.

Audit: the reconstructed one-loop decision tree

First establish the regulated self-adjoint family and its low cluster. Next compute the off-diagonal witness. Separately calculate eigenvalue motion and the required window. Then classify allowed counterterms and carrier identifications. Finally test the regulator limit. A claim of physical range motion belongs only after all five layers have been stated.

F BRST–projection equivalence: proof and limits

Let $Z^n = \operatorname{Ker} Q_n$ and $B^n = \operatorname{Ran} Q_{n-1}$. If B^n and $\operatorname{Ran} Q_n^\dagger$ are closed, then

$$\mathcal{K}^n = B^n \oplus (B^n)^\perp, \qquad Z^n = B^n \oplus \left(Z^n \cap (B^n)^\perp\right). \tag{128}$$

But

$$Z^n \cap (B^n)^\perp = \operatorname{Ker} Q_n \cap \operatorname{Ker} Q_{n-1}^\dagger = \operatorname{Ker} \Delta_n. \tag{129}$$

The quotient map $Z^n \to Z^n/B^n$ therefore restricts to a unitary identification of harmonic representatives with reduced cohomology. If $\operatorname{Spec} \Delta_n \cap (0, \gamma) = \varnothing$, the spectral theorem gives

$$\chi_{[0,\delta^2]}(\Delta_n) = \chi_{\{0\}}(\Delta_n) \tag{130}$$

for $\delta^2 < \gamma$.

The conclusion weakens in a controlled sequence:

$$\text{cohomology} \longleftrightarrow \text{harmonic kernel} \longleftrightarrow \text{zero projector} \longleftrightarrow \text{finite-window projector}. \tag{131}$$

The first arrow needs a Hodge decomposition, the second is definitional, and the third needs a gap. Comparison with a distinct master operator needs a further intertwiner. Writing the arrows separately prevents a proof of one from being cited as a proof of all.

F.1 Ordinary and reduced cohomology

If $B^n = \operatorname{Ran} Q_{n-1}$ is not closed, the algebraic quotient

$$H^n(Q) = Z^n/B^n \tag{132}$$

need not be Hausdorff in the quotient topology inherited from the Hilbert space. The reduced cohomology is

$$\overline{H}^n(Q) = Z^n/\overline{B^n}. \tag{133}$$

It is the reduced quotient that is naturally represented by vectors orthogonal to $\overline{B^n}$. If the range is closed, ordinary and reduced cohomology agree and the familiar harmonic picture follows. If the range is not closed, the difference is analytic information that a purely algebraic cohomology notation can conceal.

The same nonclosed-range phenomenon appears spectrally as the possibility that zero is not isolated in Δ_n. Approximate harmonic vectors can have arbitrarily small Laplacian norm without converging to a nonzero vector in $\operatorname{Ker} \Delta_n$. This is exactly the regime in which Klauder's nonzero window can carry information not captured by the exact zero projector.

F.2 The spectral window of the BRST Laplacian

When

$$\operatorname{Spec}\Delta_n \cap (0,\gamma) = \varnothing, \tag{134}$$

one has

$$\chi_{[0,\delta^2]}(\Delta_n) = \chi_{\{0\}}(\Delta_n) \tag{135}$$

for every $0 < \delta^2 < \gamma$. In that regime the BRST harmonic space is literally a Klauder-type spectral range, with Δ_n serving as the positive master operator.

If no gap exists, the finite-window range generally contains approximate harmonic states in addition to exact cohomology representatives. Whether those states should be treated as physical depends upon the limiting problem. One cannot infer exact equality merely because both constructions use a positive operator near zero.

F.3 Comparison with a distinct Klauder operator

Let

$$P_K = \chi_{[0,\delta_K^2]}(C_K) \tag{136}$$

act on $\mathcal{H}_K$, and let

$$P_B = \chi_{\{0\}}(\Delta_{\mathrm{BRST}}) \tag{137}$$

act on the relevant ghost-number sector of $\mathcal{K}_{\mathrm{BRST}}$. A unitary

$$U : \operatorname{Ran} P_K \longrightarrow \operatorname{Ran} P_B \tag{138}$$

identifies the Hilbert spaces. A physical equivalence requires more.

For every admitted observable A, one needs

$$U A_K^{\mathrm{phys}} U^{-1} = A_B^{\mathrm{phys}}. \tag{139}$$

For dynamics,

$$U H_K^{\mathrm{phys}} U^{-1} = H_B^{\mathrm{phys}}. \tag{140}$$

If the constructions vary over a base B, a family U_b should also satisfy

$$U_b \nabla_b^K = \nabla_b^B U_b. \tag{141}$$

Without (141), two methods can agree fiber by fiber and disagree about the result of carrying a state around a loop in parameter space.

F.4 A sufficient bundle-equivalence statement

Suppose over a connected parameter region:

1. the Klauder and BRST physical ranges are defined by persistent isolated spectral clusters;
2. their ranks or Hilbert-module types agree;
3. a continuous family of unitaries U_b identifies the fibers;
4. U_b intertwines the physical observable representations and Hamiltonians;
5. U_b intertwines the chosen physical connections;
6. the global holonomies agree under the same identification.

Then the two procedures define equivalent physical Hilbert bundles over that region. Dropping any item weakens the conclusion. State counting alone proves only state counting.

F.5 Failure modes

The comparison can fail at several layers.

Algebraic failure. The proposed BRST charge may fail to be nilpotent because of an anomaly: $Q^2 \neq 0$. The cohomology is then not defined in the stated form.

Analytic failure. The ranges of the differentials may not be closed, so harmonic representatives realize only reduced cohomology or fail to provide the intended Hilbert quotient.

Spectral failure. Zero may not be isolated. Exact harmonic projection and finite-window projection then differ.

Constraint-type failure. Second-class or partially anomalous constraints may require conversion or an extended BRST construction before a nilpotent charge exists.

Representation failure. The two constructions may land in inequivalent representations of the physical observable algebra even when abstract Hilbert spaces happen to be isomorphic.

Dynamical failure. Static state spaces may agree while effective Hamiltonians or time-dependent transport differ.

Global failure. Local gauge-fixed or BRST descriptions may glue with nontrivial transition data. One global cohomology group need not display all local-to-global comparison defects.

Window failure. The Klauder construction may deliberately retain a nonzero low band in a regime where BRST selects only an exact kernel. This is not an algebraic contradiction; the two methods are answering differently windowed questions.

F.6 Ghosts and the local–global intuition

My historical intuition was that ghosts arose from a conflict between local and global considerations. That sentence is too strong if taken as a general origin claim. Ghosts also exponentiate determinants, encode gauge directions, and organize homological cancellation on contractible perturbative patches.

The intuition retains a narrower value. A ghost complex packages local gauge redundancy into a graded algebraic construction. When local gauge choices do not glue globally, the transition and descent data become part of the story. In that setting, one global cohomological answer can obscure how positive physical representatives vary from patch to patch unless the gluing is made explicit.

The fair conclusion is not "ghosts are caused by global obstruction." It is that the BRST complex becomes especially delicate when local gauge descriptions do not assemble into a globally trivial physical family.

BRST comparison rule

Prove the arrows one at a time: cohomology to harmonic representatives; harmonic representatives to a zero spectral projection; zero projection to a finite window; BRST window to the Klauder window; and fiberwise equivalence to transport equivalence. No arrow is licensed merely because the endpoints look similar.

G Open reconstruction ledger

Question	Present status	What would close it
Which algebraic noncovariant gauge was used?	Open	Surviving notes, correspondence, code, or an independent calculation matching remembered intermediate expressions.
What was the regulated Yang–Mills master operator?	Open	Explicit smearing, ordering, domain, regulator, and counterterm prescription.
Did the one-loop correction rotate the physical range?	Open	A nonzero renormalized $(1-P_0)C_1P_0$ with a controlled gap or a continuous-spectrum replacement.
Was the apparent correction only window renormalization?	Open	Decomposition of C_1 into diagonal calibration terms and genuine off-diagonal mixing.
Does the regulator family carry curvature?	Open	At least two controlled parameters and a nonzero limit of $P(\mathrm{d}P \wedge \mathrm{d}P)P$.
When does local BRST data glue to the same physical family?	Open	Fiberwise Hodge equivalence plus compatible intertwiners, observables, states, dynamics, and trivialized coherence defects.
Can irregular kernel scalings be organized functorially?	Open	A comparison theory that preserves direct-sum sectors across window and regulator changes.

Go/no-go conclusion

The historical and mathematical essay is a go even if the lost Yang–Mills calculation is never recovered. The claim that a projection can move is certified by the 2005 model. The field-theory episode is presented as a remembered research frontier and converted into a falsifiable reconstruction program. No stronger historical claim is needed.

G.1 A reconstruction protocol

If the Yang–Mills calculation is attempted again, the first task is not to reproduce remembered intermediate expressions. It is to reconstruct the problem specification.

1. Fix the precise algebraic noncovariant gauge and residual gauge conditions.
2. Specify spatial volume, boundary conditions, ultraviolet regulator, and smearing conventions.
3. Define the regulated canonical variables and their common dense domain.
4. Write the regulated Gauss operators, including ordering choices.
5. Construct the positive closed master form and its associated self-adjoint operator.
6. Define the unperturbed physical window and establish whatever spectral separation is actually available.
7. Compute the loop correction before and after the admitted counterterms.
8. Separate the diagonal cluster correction from the off-diagonal witness.
9. Determine the running, if any, of the physical spectral window.
10. Compare regulators and gauges using explicit intertwiners rather than equal-looking formulas.
11. Test the continuum limit with the strongest convergence notion that the estimates support.

Only after these steps should the calculation be compared with remembered features of the lost notebook. Agreement with a memory can guide a search; it cannot validate an otherwise uncontrolled operator.

G.2 Possible outcomes

The reconstruction need not reproduce the interpretation I now favor. At least five outcomes would be informative.

First, the off-diagonal witness may vanish after the correct counterterms. In that case the remembered “motion” may have been only window or eigenvalue renormalization.

Second, the witness may be nonzero at finite regulator and vanish in the continuum limit. The movement would then be a regulator artifact, though possibly a useful diagnostic of how the regulated physical sectors are embedded.

Third, the witness may converge to a nonzero operator or form. That would certify genuine perturbative motion of the physical sector under the stated construction.

Fourth, the witness may diverge while a renormalized comparison or reduced kernel converges. The correct continuum object would then not be a norm-limit projector on the original carrier.

Fifth, the spectral gap may close so severely that the projector derivative is not the correct observable. Spectral flow, form convergence, or a correspondence between different carrier representations may replace the smooth Grassmannian picture.

None of these outcomes makes the reconstruction a failure. The failure would be to report "the projection moves" without saying which of them occurred.

G.3 The renormalization curvature question

At finite regulator and away from crossings, the family

$$(g, \Lambda) \longmapsto P_{g,\Lambda} \tag{142}$$

may define a smooth two-parameter physical bundle. Its mixed curvature is

$$F_{g\Lambda} = P[\partial_g P, \partial_\Lambda P]P. \tag{143}$$

If this quantity is nonzero, changing coupling and changing regulator scale do not commute as identifications of the physical sector. This does not by itself prove a physical anomaly. It says that the comparison around a small rectangle in parameter space has nontrivial holonomy to leading order.

The old attraction of the renormalization group becomes understandable in this language. The calculation had produced not only a scale-dependent operator but a potentially scale-dependent physical subspace. What was missing was the connection that tells us how to compare those subspaces.

The continuum question is harder. One must decide whether the curvature has a regulator-independent limit, converges only after rescaling, or belongs to a different carrier description. A finite-cutoff curvature is evidence for a geometric effect in the regulator family, not automatically an observable of the continuum theory.

G.4 Why the note survives a negative reconstruction

Suppose a complete modern recalculation shows that the Yang–Mills physical range does not move at the remembered order. The central mathematics of this note remains intact. The finite anomalous model still proves that projector motion is possible. The perturbation theorem still supplies the correct witness. The BRST comparison still identifies its hypotheses. The reduced kernel theorem still explains singular survival.

The historical conclusion would become more precise: an early calculation created the impression of motion, and that impression helped expose the right geometric distinction even though the particular field-theoretic witness vanished. That would be a legitimate scientific history. A question can be fruitful without its first suspected answer being correct.

Conversely, if the witness survives, the note will already contain the language needed to state the result without exaggerating it. Either way, the reconstruction is protected from becoming a legend whose value depends upon a single unrecoverable outcome.

Audit: hostile-reader test

A skeptical reader should be able to delete every remembered claim and still find a coherent mathematical note: positive constraint compression, reduced kernels, a certified finite moving projector, spectral perturbation, BRST comparison, and the geometry of physical ranges. The memories explain why these pieces belong together; they do not bear the weight of their proofs.

H Methodological postscript: reconstruction without possession

H.1 Five evidentiary postures

The five labels used in this note correspond to five different relations between the present author and the claim being made.

Recorded means that a surviving object can be inspected by another reader. A publication, dissertation, dated manuscript, or correspondence can support a claim independently of my present memory. Recorded evidence can still contain an error, but the historical existence of the statement is not dependent upon recollection.

Remembered means that the claim belongs to autobiographical testimony. Such testimony is appropriate for conversations, motivations, emotional responses, and the order in which problems entered my attention. It is not sufficient to establish a technical expression that no longer survives.

Reconstructed means that the mathematics has been derived again from assumptions now stated. A reconstruction can certify that the mechanism works. It cannot by itself prove that every step occurred in precisely that form at the earlier date.

Modern reading means that later mathematics has supplied a language or organizational principle that I did not possess in the original period. The Grassmannian speed, Kato connection, and mixed renormalization curvature belong to this posture as used here. They illuminate the old question without being projected backward as remembered notation.

Open means that the note has reached the evidentiary boundary. An open claim is not a failure of prose waiting to be filled. It is an explicit research obligation.

H.2 Reconstruction is not replication

Replication asks whether a specified calculation or experiment produces the same result again. Reconstruction begins earlier: what exactly was the calculation? When the gauge choice, regulator, ordering, or counterterm scheme is missing, several mathematically reasonable replications can disagree without any of them reproducing the historical object.

The first goal is therefore not numerical agreement with memory. It is to build a well-posed family of operators that expresses the remembered question. Only after the family is controlled can one ask whether its behavior resembles the original report.

This is why the off-diagonal witness is more useful than a remembered phrase. "The projection needed a correction" admits several meanings. The equation

$$(1 - P_0)C_1P_0 \neq 0 \tag{144}$$

specifies one. A new calculation can establish or refute it.

H.3 What a language model contributed

The present expansion was assisted by a language model as a tool for staging, bookkeeping, comparison, and exposition. The mathematical direction came from the remembered problem, the surviving papers, and the later framework developed by the author. The tool helped keep the evidentiary categories separate, formulate the perturbative alternatives, and turn a compressed twenty-page draft into a note with room for the intellectual history.

That division of labor matters. A model can generate a plausible bridge between two equations. The bridge becomes mathematics only when its hypotheses and proof can be checked. It can generate a plausible historical narrative. The narrative becomes evidence only where a source or labeled memory supports it.

The tool is especially good at removing the silence around an unfinished question. The author remains responsible for deciding when silence is the correct answer. In this note, "Open" is sometimes the most important word on the page.

H.4 Why preserve the discovery path?

A formal paper usually removes the order in which the ideas were found. The definitions appear before the theorem; the failed interpretation disappears; the decisive example looks inevitable. That organization is useful for verification. It is incomplete as an account of inquiry.

Here the order matters. Reduced kernels came before the modern language of varying carriers. The anomalous model proved motion before I knew that motion was the relevant conclusion. The Yang–Mills calculation created discomfort before projector perturbation supplied the witness. The dissertation advice separated what could be completed from what had to wait. The later fiber view did not generate this history; it made the history intelligible.

Preserving that order lets a reader see how mathematical taste develops. One learns not only the final criterion but why a physicist might fail to ask for it at first. That knowledge can be more reusable than another polished derivation whose motivating obstruction has been erased.

H.5 The right to leave a question open

An era of rapid formal generation can make an unanswered question feel like a defect in presentation. But a forced answer damages both mathematics and memory. The right to leave the Yang–Mills witness open is what permits the rest of the note to be strong.

If the calculation is reconstructed, the ledger can be updated. If it is not, the note still records the exact experiment that would decide the issue. An open question with a witness, a domain, and a failure criterion is not empty. It is a location on the map.

This is the broader purpose of the View from the Mountain series. The mountain view can connect paths that looked separate below. It can show that a small calculation planted a durable seed. It can also show which ridge was never crossed. Accuracy requires both kinds of recognition.

Final reconstruction rule

Use memory to recover the question. Use surviving work to locate the old mathematics. Use present mathematics to reconstruct only what its assumptions support. Label the modern interpretation. Leave the remaining ridge open.

References

[1] A. Bassetto, G. Nardelli, and R. Soldati, *Yang–Mills Theories in Algebraic Non-Covariant Gauges: Canonical Quantization and Renormalization*, World Scientific, Singapore (1991).

[2] W. A. Bomstad and J. R. Klauder, "Linearized quantum gravity using the projection operator formalism," arXiv:gr-qc/0601087 (2006).

[3] M. Henneaux and C. Teitelboim, *Quantization of Gauge Systems*, Princeton University Press (1992).

[4] T. Kato, *Perturbation Theory for Linear Operators*, corrected printing of the second edition, Springer (1995).

[5] J. R. Klauder, "Universal procedure for enforcing quantum constraints," *Nuclear Physics B* **547**, 397–412 (1999).

[6] J. R. Klauder and J. S. Little, "Elementary model of constraint quantization with an anomaly," *Physical Review D* **71**, 085014 (2005), arXiv:gr-qc/0502045.

[7] J. R. Klauder and J. S. Little, "Highly irregular quantum constraints," arXiv:gr-qc/0602114 (2006).

[8] N. Aronszajn, "Theory of reproducing kernels," *Transactions of the American Mathematical Society* **68**, 337–404 (1950).

[9] M. Reed and B. Simon, *Methods of Modern Mathematical Physics, Vol. I: Functional Analysis*, Academic Press (1980).

About *A View from the Mountain*

A View from the Mountain is a series of historical and mathematical research notes by J. Scott Little. Each note returns to an earlier line of inquiry, reconstructs what can still be certified, records what was learned from the failed or unfinished portions, and compares the original view with the wider geometry visible from later work.

The series uses five evidentiary labels throughout:

Recorded	supported by a surviving source;
Remembered	autobiographical chronology or motivation;
Reconstructed	mathematics rederived from stated assumptions;
Modern reading	a later interpretation not attributed backward;
Open	a ridge that the present evidence does not cross.

The purpose is not to make the past look inevitable. It is to preserve the progression of thought while allowing proof, source, and explicit uncertainty to determine what may be claimed.

Research Note No. 03

The Projection Operator Moves

Constraints, reduced kernels, BRST, and the geometry of physical ranges. The finite mechanism is certified; the lost Yang–Mills one-loop witness remains open. Every mountain needs a goat. This one had John Klauder.