Enacting Rigour:

Global-Rebuilding and Local-Repair in Quantum Field Theory

James D. Fraser[*]

Forthcoming in *Philosophy of Rigour*, D. Rickles and K. P. Y. Thébault (eds), Routledge.
Please cite published version.

Abstract

Quantum field theory (QFT) has long been subject to rigour concerns. Attempts to reformulate the theory in a more rigorous form have not converged on a single formalism, however, but have instead produced a plurality of distinct approaches. This chapter treats these reformulation programmes as a case study for investigating the methodological role of rigour in science. Tracing the development of axiomatic QFT and causal perturbation theory, I argue that these programmes instantiate distinct rigourisation strategies: a global-rebuilding strategy, which seeks to reconstruct the target theory starting from a new mathematical foundation, and a local-repair strategy, which aims to resolve specific points of unclarity within the existing formalism. The philosophical upshot is that we need to conceive of rigour as an ideal that is enacted differently on different methodological pathways. I briefly touch on how this framework might be extended to other areas of mathematics, science, and philosophy where rigour concerns loom large.

1 Introduction

The ambition to make quantum field theory (QFT) rigorous has not produced a single formalism; instead, it has given rise to a plurality of reformulations of the theory. Besides well known programmes like axiomatic QFT we also find independent traditions, like causal perturbation theory, that pursue rigour in markedly different ways.

How can we account for this diversity? This is the question broached in this chapter. In doing so, I engage both with the local question of how to negotiate the plurality of QFT formalisms and the more general question of how rigour figures in the methodology of science.

Within the philosophy of physics literature, the plurality of QFT formalisms has been most visibly discussed in the debate between Doreen Fraser and David Wallace [Fraser,

---

[*] Aix-Marseille University. Email: jamesf09@hotmail.co.uk

2009, 2011, Wallace, 2006, 2011]. The present chapter is not another entry in that debate. Rather than offering reasons to select one formalism as the correct object of philosophical interpretation, the aim here is to make sense of how such multiplicity arose in the first place. This is a historical question, but is also a philosophical one, since it requires a model of the methodological decisions that shaped each programme. Historically, the chapter contrasts the development of axiomatic QFT and causal perturbation theory from the 1950s to the 1970s. Philosophically, it treats rigour not as a selection criterion but as a lens for reconstructing the methodological branch points that produced the diversity of formulations we see today.

A disclaimer about the limited scope of my discussion is in order. Ultimately, what is needed is a more expansive and fine-grained map of the full space of QFT formulations. The present discussion makes a modest contribution towards this larger goal by charting one sub-region of that space: the space of rigour-motivated reformulations of QFT that emerged in the formative period of mathematical physics. To complete the picture, one would also need to analyse the arguably considerably more diverse space of formulation programmes sometimes lumped under the heading of "conventional" QFT, which is not attempted here.

While the plurality of rigour-orientated approaches to QFT has some intrinsic interest to philosophers of physics, a more tangible payoff of my analysis is that it provides a rich context to ask more general philosophical questions about the methodological role of rigour. Much of the emerging philosophical literature on this subject asks the quintessential conceptual analysis question, what is rigour?[1] By contrast, I will not have much to say about what rigour is here. Indeed, I shall take for granted a thin practice-oriented conception of rigour as a normative standard governing clarity, control, and explicitness in definitions, constructions, and arguments. My focus is on what rigour does rather than what it is; on how demands for rigour are motivated and implemented in particular scientific contexts.

The key claim is the following: making a branch of physics rigorous is not a mechanical operation that takes in a messy formalism and spits out a pristine one. Rather, rigourisation takes place through a broader strategic response to local scientific problems. This means that there can be multiple ways to enact rigour in a given context. In the case of QFT I distinguish two rigourisation strategies:

1. A *global-rebuilding* strategy: which tries to recover the successes of an existing theory from a new set of objects and conditions whose credentials are under control from the start. I associate this strategy with axiomatic QFT.

2. A *local-repair* strategy: which takes an existing theory in its current state and addresses specific points of unclarity that arise within it. I associate this strategy

---

[1] Tanswell [2024] and Ashton [2024] are recent reviews which collate and categorise the various views of the nature of mathematical rigour which have been put forward in the philosophy of mathematics literature.

with causal perturbation theory.

Distinguishing these strategies gives us resources to explain the plurality of rigour-orientated reformulation programmes. It also gives us a language to describe how they evolved over time: a key sub-claim of my discussion is that the global-rebuilding and local-repair strategies are themselves not monolithic, but are adapted to the perceived risks and opportunities of a given state of theoretical development. An open question which is briefly touched on at the conclusion of the chapter, and could be explored further, is the extent to which this model of the methodological choices involved in rigourisation can be usefully applied to other areas of science and mathematics.

This chapter is structured as follows. Section 2 sets the scene by explaining the factors that led rigour to become an operative ideal in 1950s high-energy theory. Sections 3 and 4 cover the axiomatic QFT and causal perturbation theory traditions, respectively, presenting them as diverging strategic responses to this original problem situation. Section 5 discusses how the rise of quantum chromodynamics (QCD) in the 1970s altered the problem situation; this gives an opportunity to emphasize the flexibility of global-rebuilding and local-repair strategies. Section 6 concludes with some more general reflections on the distinction between global-rebuilding and local-repair.

2 1950s: The Original Problem Situation

Axiomatic QFT and causal perturbation theory both originated in the 1950s. The aim of this section is to characterise what I call the *problem situation* from which these rigour-orientated reformulation programmes emerged.[2] By a problem situation, I mean a set of problems together with evaluative standards relative to which those problems are identified, prioritised, and assessed. Importantly, this can include uncertainty about the nature of the problems themselves and about what would count as an adequate response to them. It is against this background that a methodological strategy is formulated.

In the case of 1950s QFT, the problem situation was characterised by an accumulation of *mysterious internal anomalies*. By internal anomalies, I mean problems that arise within the formalism itself, as opposed to a mismatch between formalism and empirical data. Physicists in the 1950s and 1960s frequently described these issues as threats to the consistency of QFT. Yet, as Blum [2026, Forthcoming] documents in detail, the problems in question were heterogeneous and rarely took the form of straightforward logical contradictions. More commonly, they appeared as various kinds of mathematical

---

[2] I borrow the term "problem situation" from Popper [1972]. He characterises it, loosely as I have here, as the background context against which a methodological strategy is selected and rationalised. One might question whether there really is a shared conception of this background. Seth [2007]'s push back against attributions of a universally perceived sense of "crisis" is relevant here. My abrupt periodisation of the 1950s and 1970s problem situations is also highly idealised. While the methodological model sketched in this chapter is somewhat naive, I hope that it is a useful starting point which can be sophisticated through more detailed engagement with the historical record.

divergence.

What is crucial for my purposes is not merely the presence of such anomalies, but their mysteriousness. It was unclear whether these divergences signalled a genuine breakdown of QFT, or whether they instead reflected inadequacies in its mathematical implementation. It was this diagnostic opacity—rather than the mere existence of technical difficulties—which fueled the imperative to improve standards of rigour.

In what follows, I give a selective survey of some of the most important internal anomalies that motivated the drive towards increased rigour.[3]

The earliest worries about the consistency of QFT go right back to the late 1920s and the discovery of divergent integrals in QFT perturbative expansions. Heisenberg and Pauli [1929]'s early formulation of quantum electrodynamics (QED), already hit upon this problem. They treated quantum fields as operator-valued functions of space-time, with associated (equal-time) commutation rules, and constructed the theory's Hamiltonian from products of these fields. In order to describe interacting field theories (whose nonlinearity made exact solutions unreachable) they proposed expanding quantities in a perturbation series. The Hamiltonian of the theory is split into a "free" part and an interaction term:

$$H = H_0 + \lambda H_I, \tag{1}$$

with the observables of the full theory then being expressed as power series in the coupling λ.

Quite unexpectedly, however, it was found that the coefficients of this series contain integrals that diverged due to the contribution of arbitrarily high momentum modes of the fields. This is the problem of ultraviolet divergences.

On the face of it, this is not just an internal anomaly but an empirical catastrophe, since it seemed to imply that the theory ascribed infinite values to measurably finite quantities. Diagnoses of the root of the problem varied, however. Some theorists took the divergences to signal an inconsistency between the physical principles of QFT, so that the theory had to be abandoned, or at least fundamentally modified. The most important anti-QFT programme for our story was the S-matrix theory which Heisenberg initiated in the last years of the second world war [Heisenberg, 1944, 1943a,b]. Heisenberg proposed starting from a new object—the S-matrix, an operator mapping incoming asymptotic states at $t \to -\infty$ to outgoing states at $t \to \infty$. He then tried to build up a new approach to relativistic quantum theory by imposing conditions such as unitarity and Lorentz invariance directly on this asymptotic evolution operator.

A less drastic response was pursued by the architects of perturbative renormalisation at the end of the 1940s. Rather than rejecting perturbative QFT tout court, theorists like Feynman, Schwinger, Tomonaga and Dyson, developed a way of handling the infinities that appeared in Pauli and Heisenberg's original formalism. In this framework, a new set

---

[3] Complementary discussions which cover a broader range of internal anomalies, or give more details about those listed here, are Blum [2026, Forthcoming], Blum and Fraser [2025], Fraser and Rejzner [2024].

of (finite) renormalised masses and interaction couplings are introduced, with the original (infinite) “bare” masses and couplings being cancelled by (also infinite) counterterms. In the case of theories like QED it proved to be possible to eliminate all of the ultraviolet divergences in the perturbation series in this way via the redefinition of a finite number of dynamical parameters.

While the first few terms of these renormalised perturbation series now yielded extremely accurate predictions, this did not resolve the sense of unease surrounding the foundations of QFT. One issue was that the manipulations which went into the renormalisation procedure appeared to many to be mathematically unsound. The "bare parameters", invoked in the procedure, were identified with power series with infinite coefficients, and thus appeared to be ill-formed expressions. This created a suspicion that the renormalisation procedure had papered over the pathologies signalled by the ultraviolet divergences rather than genuinely addressing them. Given these shaky mathematical foundations, the stupendous success of renormalised perturbation theory itself presented an explanatory problem for foundationally minded theorists: why does this dubious method work so well?[4]

Indeed, despite newfound empirical success, the 1950s saw an intensifying of foundational worries about QFT, as a new set of internal anomalies started to pile up.

One disturbing realisation was that, even after the divergences in the series coefficients had been handled via renormalisation, the large-order behaviour of the QFT perturbation series was probably divergent. Freeman Dyson had hoped to prove the existence of exact solutions of QED by establishing the convergence of the theory’s renormalised perturbation series, but ended up developing an argument for the opposite conclusion Dyson [1952]. This verdict was further bolstered by studies of toy QFT models that displayed divergent large-order behaviour [Thirring, 1953, Hurst, 1952]. While this did not indicate inconsistency on its own, it did point to fundamental limitations of the perturbative treatment of interactions. Even if the internal issues with the perturbative expansion could be resolved, it seemed that the perturbation series on its own was incapable of establishing the existence or otherwise of exact solutions. A cavern had thus opened up between the empirically successful perturbative framework and the theory’s exact, now dubbed non-perturbative, content.[5]

In fact, the situation was even more fraught, as the formulation of dynamical equations about which we might ask about the existence of solutions was itself problematised. Dyson’s canonical derivation of the perturbative expansion had proceeded by integrating

---

[4] In other fields of pure and applied mathematics, we also see inexplicable successes, in addition to paradoxes and inconsistencies, as drivers of rigourisation. Potential examples would be the success of infinitesimals in calculus and the quantitative accuracy of divergent series expansions like Stirling’s approximation. In both cases, later rigour-orientated reformulation programmes were given impetus by the undeniable success of these apparently problematic techniques. This suggests a distinction between positive and negative motivations for rigourisation, which is worthy of further exploration.

[5] For more discussion of the divergence of QFT perturbation theory in historical context, see Blum [2026] and Fraser and Rejzner [2024].

the so-called Schwinger-Tomonaga equation: a relativistic version of the Schrodinger equation which evolves states between space-like hypersurfaces using the interaction Hamiltonian only. However, Haag [1955] argued that, under natural assumptions, interacting fields could not be represented on the same Hilbert space as free fields. This apparently implied that the unitary operator $U(t, t')$ which Dyson obtained by integrating the Schwinger-Tomonaga equation could not, in fact, exist. A less well-known episode that is significant for our discussion later on, is that Stueckelberg [1951] discovered yet another type of divergence (so-called boundary divergences) when trying to calculation Dyson's $U(t, t')$ operator for finite time arguments. This contributed to the sense that the dynamical framework assumed in the perturbative formulation of QFT could not be transferred to a non-perturbative context, and some other approach to dynamics was needed.[6]

Finally, and most alarmingly of all, new non-perturbative arguments for the inconsistency of QED were put forward by theorists like Källén and Landau in the mid-1950s [Källén, 1953, Landau and Pomeranchuk, 1955, Landau et al., 1956]. These arguments hinged on writing down expressions for the renormalised charge of QED outside of perturbation theory and then trying to demonstrate that it displayed divergent behaviour. Landau, in particular, gave his name to the so-called Landau pole: an infinity in the renormalised charge at some very large, but finite, energy scale. While they targeted QED specifically, these arguments were widely taken to threaten the existence of solutions to any interacting QFT, independently of the status of the renormalised perturbative expansion.

By the late 1950s, quantum field theory thus faced a complex, layered predicament: perturbation theory appeared incapable of answering consistency questions; the standard dynamical frameworks that might be used to pose such questions were themselves under strain; and independent arguments suggested that no satisfactory interacting theory existed at all.

Crucially, however, there was still no clear-cut demonstration of the inconsistency of the QFT framework. If such a thing existed the methodological upshot would also be clear-cut: QFT must be rejected and replaced by a new theory. In this scenario, there would be no imperative to improve standards of mathematical rigour.

The crisis facing QFT in the 1950s, however, was stubbornly mysterious and open-ended, leaving a deep sense of uncertainty about the developmental direction of fundamental physics. As Streater and Wightman [1964] put it:

> ...[T]he Main Problem of quantum field theory turned out to be to kill it or cure it: either to show that the idealizations involved in the fundamental notions of the theory (relativistic invariance, quantum mechanics, local fields, etc.) are incompatible in some physical sense, or to recast the theory in such

---

[6] A much more detailed discussion of these problems with the Schwinger-Tomonaga equation, and how this leads into the story of causal perturbation theory, can be found in Blum and Fraser [2025].

> a form that it provides a practical language for the description of elementary particle dynamics. [Streater and Wightman, 1964, 1]

It was the diagnostic opacity of the internal anomalies, I suggest, more than anything, that drove attempts to reformulate QFT in more rigorous terms.

The central claim of this chapter is that axiomatic QFT and causal perturbation theory can be understood as embodying different strategic responses to this problem situation—specifically, different ways of dealing with diagnostic opacity. The local-repair strategy, exemplified by causal perturbation theory, deploys rigour in the service of internal clarification: it aims to demystify individual anomalies by subjecting them to progressively finer-grained mathematical analysis. The global-rebuilding strategy, exemplified by axiomatic QFT, adopts a markedly different stance. Rather than attempting to diagnose the pathologies of the inherited formalism, it brackets them and constructs a new, ambiguity-free framework in their place. If this reconstruction succeeds, the original mysteries are not so much resolved as rendered redundant.

The following two sections flesh out this picture by examining the key features of axiomatic QFT and causal perturbation theory, respectively, in their formative years.

3 Axiomatic QFT

Axiomatic QFT is a complex, multi-stranded movement, and its intellectual history remains largely unwritten (Blum [2026] is an invaluable starting point). I will not attempt this herculean task here. Rather, I touch on episodes from the development of the tradition that are relevant to diagnosing its methodological rationale—or, more accurately, its evolving rationales (plural). In particular, this section presents axiomatic QFT as one possible response to the problem situation posed by the cluster of internal anomalies described in the previous section.

A key early decision that shaped the axiomatic QFT programme was its rejection of renormalised perturbation theory. From the outset, the ambition was not to shore up the existing perturbative formalism, but to develop a novel and entirely non-perturbative framework for quantum field theory.

This can be understood as a reasonable—though by no means the only—response to the problem situation of the 1950s. As we saw, renormalised perturbation theory had acquired a dense knot of internal difficulties, and there were so many sites of mathematical unclarity that it was difficult to identify the fundamental source of the trouble. Moreover, even if these difficulties could be resolved—if the knot could somehow be untied—there were good reasons to think that this would not settle the most basic questions about the consistency of the QFT framework. As Arthur Wightman, a key architect of axiomatic QFT, later put it in a retrospective account of the period:

> "[A]re the renormalized series convergent or divergent? Although there is even

> yet no proof in quantum electrodynamics, it is generally believed on the basis of a study of simpler cases that they are divergent... Thus at this stage in the development, which occurred in the early 1950s, the key problem was to find some non-perturbative approach to the solutions of quantum field theory." [Wightman, 1976, 195]

For Wightman then, answering the core foundational questions about QFT meant moving away from perturbation theory. Given that interacting QFTs had so far been described perturbatively, this essentially meant clearing the desk and beginning again.

From a rigourisation perspective, this had a clear advantage. Mathematical control could be established at the outset and then maintained as the new formalism was elaborated. Thus, axiomatic QFT exemplifies what I am calling a global-rebuilding approach to rigourisation: rather than diagnosing and repairing problems in the existing formalism, it aims to recover its successes starting from a novel, and demonstrably secure, foundation. If this can be done, then the internal problems with QFT would all be resolved at once, as it were.

In what sense, exactly, was this project "axiomatic"? This was not an attempt to recast QFT in logical or set theoretic terms for the sake of meta-mathematical clarification—"formal axiomatics" in the terminology of Rédei and Stöltzner [2006]. Rather, this was "soft axiomatics", with postulates being stated in terms of higher-mathematical structures and conclusions being drawn from them using ordinary mathematical arguments. Relevant archetypes were von Neumann's axiomatization of quantum mechanics and Hilbert's work on the axiomatization of classical physics. However, I will suggest below that the methodological context in which early axiomatic QFT operated was quite different from these predecessors, and that assimilating it straightforwardly to Hilbert's 6th problem obscures some of the programme's distinctive features.

It is useful to distinguish three parts of the axiomatic QFT style reconstruction programme:

1. A class of basic objects is specified.
2. A set of mathematically precise conditions, the axioms, are imposed on these objects.
3. There were two main avenues for elaboration: the construction of models satisfying the axioms, and the derivation of general, model-independent theorems.

Let us consider each of these steps in turn.

The first step, of defining basic objects, was actually a highly non-trivial one, and turned out to be the juncture at which different sub-traditions of axiomatic QFT diverged from one another.

One reason this point was so crucial was that many of the internal issues with QFT can be traced back to its naive basic objects—quantum fields. Heisenberg and Pauli

[1929] had treated quantum fields as assignments of operators to points of space-time. Heisenberg's war-years intuition that the localisation of field operators to points was at the root of many of the theory's problems turned out to be surprisingly prescient. Rather than adopting Heisenberg's drastic solution of abandoning local objects entirely and working with asymptotic S-matrices, however, axiomatic QFT practitioners generally worked to rehabilitate local objects.

Doing so required importing new mathematical machinery. For Wightman and his collaborators, distribution theory was the key resource. Circa 1950, the idea that the classical notion of a function could be generalised to include more singular objects had only recently become well established in pure mathematics.[7] Identifying distributions with maps from an appropriately constrained space of test functions to numbers allowed objects like the Dirac delta "function", which cannot be understood as maps between number fields, to be accommodated into a more expansive space of objects. Restricting quantum fields to space-time points (or indeed hypersurfaces) leads to divergences even within free QFTs, because they too display distributional singularities. Thus, in Wightman's framework, the basic objects were operator-valued distributions: the definition of a quantum field had thus been repaired by "smearing" them against test against functions.[8]

Later on, a second round of mathematical importation, this time from the theory of operator algebras, led to an even more revisionary answer to the basic objects question.[9] In Haag and Kastler [1964]'s axiomatic system, the basic objects were taken to be C*-algebras associated with open-bounded regions of Minkowski space-time. Quantum fields were thus dethroned as fundamental objects, being instead viewed as derived representations. The motivations which prompted this branching off from the Wightman Hilbert space-based approach deserve careful historical analysis, but we can safely list one key motivation: because QFTs have infinitely many degrees of freedom the canonical commutation relations generically admit unitarily inequivalent Hilbert space representations. Ascending to the higher level of abstraction afforded by operator algebras provided a new language for describing these structures.[10] Crucially, operator algebras were associated with open bounded regions, rather than points or surfaces, encoding the distribution theoretic insight that one needs to start with smeared rather than pointlike objects in

---

[7] Schwartz [1951]'s book was responsible for inaugurating distribution theory as a distinct branch of functional analysis. However, as I touch on later, Sobolev [1936] had earlier developed many key ideas around singular generalised functions in the more specific context of partial differential equations.

[8] While Wightman's axiomatic system apparently has quite an early origin, in the 1950s, canonical statements of the axioms were not published until Streater and Wightman [1964]. Wightman [1996] is a useful retrospective discussion of how distribution theory entered into the axiomatic QFT programme.

[9] It is perhaps more correct to say that the importation of operator algebra theory was simultaneous with distribution theory but that it took longer for algebraic ideas to catch on. Irving Segal is the key early proponent of the relevance of operator algebra theory to QFT (see Segal [1959], for instance, for a precursor of later algebraic axiomatisations). This pre-history of the algebraic approach to QFT stands in need of further historical investigation.

[10] This offered an improved diagnosis of another class of divergences which were not discussed in section 2, namely infrared divergences. See Fraser and Rejzner [2024] for some discussion of their role in this story.

QFT.

Once the basic objects had been selected, the next step was to impose global conditions upon them: the axioms. Rather than attempting an exhaustive catalogue of the axioms proposed across the tradition, it is more illuminating to distinguish the different roles such axioms played, and the different kinds of justification appealed to in motivating them.

At one end of the spectrum were axioms that were treated as effectively non-negotiable. These were conditions either taken to be well confirmed within the intended empirical domain of QFT, or regarded as constitutive of what it means to have a relativistic quantum theory at all. Poincaré covariance is the clearest example of this category. Abandoning such principles would not amount to revising QFT, but to changing the subject.

Other axioms had a more complex concoction of physical and pragmatic motivations. A widespread conviction across several strands of high-energy theory in the 1950s was that Poincaré covariance alone was insufficient to capture the local correlational structure characteristic of relativistic quantum systems, and that some additional causality condition was required (see Blum and Fraser [2025] for a broader discussion of this theme). While different traditions proposed different ways of articulating this requirement, the condition that came to play this role within axiomatic QFT was microcausality. In Wightman's framework, this was expressed as

$$[\phi(x), \phi(y)] = 0 \quad \text{if } (x-y)^2 < 0, \tag{2}$$

that is, field operators associated with spacelike separated regions commute (a condition that can be translated into the algebraic language as the condition that algebras associated with spacelike separated regions commute).

Physical considerations—in particular, a prohibition on superluminal signalling—were often invoked to support this condition. But derivational power also played a crucial role in establishing its central status. The microcausality condition actually first appears as an assumption in Pauli [1940]'s proof of the spin–statistics theorem, and it later came to play a supporting role in many structural results that were taken to characterise relativistic quantum theories. Later on, the condition was also used to derive analyticity properties that figured in ambitions towards genuinely predictive non-perturbative calculations.[11] Microcausality was viewed as an attractive axiom, therefore, because it was a minimal constraint on the dynamics which yielded a substantial increase in derivational power.

At the other end of the spectrum from Poincaré covariance, there were various conditions imposed within axiomatic frameworks that had an overtly technical character. These were introduced primarily to ensure that particular derivations went through

---

[11] Papers like Goldberger [1955] used microcausality to secure analyticity properties needed to derive dispersion relations. Attempts to make these derivational connections more rigorous were often continuous with the axiomatic QFT tradition, as is briefly touched on, in connection with the Klein [1961] volume, below.

smoothly, rather than because they were regarded as expressing deep physical principles. Examples include various domain assumptions, as well as the asymptotic conditions employed in the original LSZ formulation of scattering theory [Lehmann et al., 1955]. Such assumptions were often treated as revisable, and in some cases were modified or eliminated in later iterations of the axioms, consonant with what Stöltzner [2001] calls "opportunistic axiomatics".

These considerations lead naturally to the final question: what was to be done with an axiomatic system once its basic objects and postulates had been (provisionally) fixed? As suggested above, there were two broad routes for elaborating the theory.

The first was the construction of models satisfying the axioms. For most of the 1950s and 1960s, this was not where the main focus of attention was (as we shall discuss in section 5, constructive field theory really took off later on). However, there was one way in which model construction did play a crucial role in the formative years of the programme. Constructing free quantum field theories using the newly imported mathematical frameworks, and showing that they satisfied a minimal set of axioms plausibly taken to be constitutive of QFT, was a decisive advance. As we discussed in section 2, one of the challenges facing QFT in the 1950s was providing a clear formulation of the consistency question, since the status of putatively defining dynamical equations like the Schwinger-Tomonaga equation was itself in doubt. Early axiomatic QFT succeeded in providing a clean and minimal formulation of the consistency question and answering it in the affirmative: since models of the axioms exist, they must be jointly consistent. As Blum [2026] emphasises, this was arguably the main achievement of early axiomatic QFT.

This did not, of course, resolve the broader crisis facing QFT, since it was interactions that generated the most persistent internal anomalies. One might therefore expect the next step to have been the construction of interacting models—perhaps even QED itself. Rather than attempting explicit model construction, in the 1950s-1960s period, axiomatic QFT largely concentrated on deriving model-independent theorems from the axioms alone. These included results such as the spin–statistics theorem, the CPT theorem, and structural classification results for nets of C*-algebras in the Haag–Kastler framework.

To some extent, this model-independent approach was understandable and was indeed consonant with the broader community feeling in high energy theory at the time. There was deep uncertainty about what a satisfactory quantum field theoretic description of nuclear interactions would look like, or even whether such a description was possible. At the same time, phenomenological progress outside perturbation theory was being made via dispersion-relation techniques, which derived constraints on observable quantities directly from analyticity properties—often without reference to a specific Hamiltonian dynamics. As was mentioned earlier, microcausality functioned as a minimal assumption guaranteeing the analyticity properties required for such derivations.

From this perspective, axiomatic QFT and dispersion-relation phenomenology can be seen as forming an uneasy but mutually reinforcing alliance. The axiomatic framework offered a more stable justificatory foundation for analyticity-based reasoning, while dispersion-relation methods provided concrete phenomenological traction in a domain where explicit dynamical models were unavailable. An intriguing 1961 collection titled "Dispersion Relations and the Abstract Approach to Field Theory" brought together contributions from both traditions, testifying to the sense that these seemingly very different projects were, for a time, moving in the same direction [Klein, 1961].

Having surveyed the core features of the axiomatic QFT programme, we can now return to the question of its overarching methodological posture. A common way of framing axiomatic QFT—endorsed at times by Wightman himself in retrospective remarks [Wightman, 1976], and in the philosophical literature by Rédei [2014]—is to treat it as an application of Hilbert's Sixth Problem to quantum field theory. This reading is not mistaken. There is a clear line of influence from Hilbert's blueprint for the axiomatisation of physics and the axiomatic QFT tradition. Furthermore, Hilbert's programme is certainly an exemplar of the broader global-rebuilding rigourisation which I take axiomatic QFT to be another instantiation of.

What this characterisation risks obscuring, however, are the respects in which axiomatic QFT, especially in its formative decades, was quite different from other Hilbert inspired axiomatisation projects—the ways in which the global-rebuilding strategy was adapted to the specific problem situation of QFT in the 1950s.

In the case of classical mechanics, fluid mechanics, and even quantum mechanics, axiomatisation programmes took place against a background of high confidence in the scientific adequacy of the theory under consideration. By contrast, axiomatic QFT emerged under conditions of profound uncertainty about whether interacting quantum field theories even existed as mathematically coherent entities. In this context, the meticulous definition of basic objects and the careful articulation of axioms served not merely reductive or explanatory functions, but a stabilising one. Global rebuilding created what might be called a safe arena: a cordoned-off domain in which rigorous results could be established without presupposing the reliability of QFT broadly construed.[12]

Another way to complicate the picture is to recognise that, in many ways, Heisenberg's S-matrix programme provided a more immediate template for early axiomatic QFT than Hilbert's 6th problem. Heisenberg's proposal to abandon local field dynamics in favour of asymptotic observables exemplified a radical response to the perceived pathologies of Lagrangian quantum field theory. Although axiomatic QFT practitioners were considerably more cautious, their early work shared with Heisenberg a distrust of local evolution equations and a willingness to explore genuinely novel theoretical frameworks based on

---

[12] Alongside Rédei and Stöltzner [2006]'s and Stöltzner [2001]'s notion of "soft axiomatics" and "opportunistic axiomatics", one could consider adding "existential axiomatics" to describe this use of an axiomatic method to manage uncertainty about the very existence of models.

globally imposed conditions. Whereas Hilbert and von Neumann did not expect their axiomatic formalisms to uncover new empirical results, this was a live possibility in the case of axiomatic QFT. As I pointed out above, there were tentative connections between axiomatic QFT frameworks and non-perturbative calculational methods which many hoped would entirely displace the perturbative approximation scheme.

One thus finds in early axiomatic QFT an ambition not just to regiment and clarify an existing formalism, but to transcend it. As Wightman revealingly commented in an interview with Jagdish Mehra:

> Initially, I believe that LSZ took the view that it was a considerable advantage to work with their formalism because you didn't have to go down to the disgusting problems of Lagrangian field theory. To some extent, when you have a new formulation of things, you can celebrate that and contrast it with the old. Wightman [1987].

Alongside its characteristic existential caution, then, early axiomatic QFT could accommodate novel theory-building ambitions which are not typical of Hilbertian axiomatic reconstructions.

It was arguably only later, as attitudes to QFT shifted in the 1970s, that axiomatic QFT aligned itself with a more recognisably Hilbertian self-image, as we shall discuss further in section 5. For now, we leave axiomatic QFT and consider the other rigourisation programme of the 1950s: causal perturbation theory.

4 Causal Perturbation Theory

While axiomatic QFT has received the lion's share of philosophical attention, it is not the only rigour-orientated QFT reformulation programme. A substantial body of work has also been devoted to clarifying the mathematical structure of the original perturbative formalism, with the so-called *causal perturbation theory* (CPT) framework serving as the central exemplar of this tradition.[13] The early history of this programme has recently been reconstructed in detail by Blum and Fraser [2025]. In this section, I review some key episodes from that history, presenting causal perturbation theory as an alternative rigourisation response to the problem situation of the 1950s. In contrast to axiomatic QFT's attempt to rebuild the theory from new foundations, CPT proceeds by intervening within the inherited perturbative framework itself, seeking to diagnose and resolve specific sources of mathematical unclarity.

Causal perturbation theory diverges methodologically from axiomatic QFT from the outset by taking a diametrically opposed stance toward perturbation theory. This contrast is already visible in the differing attitudes of Wightman and Bogoliubov. Whereas

---

[13] Another strand of this literature, which would also be interesting to examine from the perspective of rigourisation, is the BPHZ formalism (also initiated by Bogoliubov).

Wightman treated perturbation theory as a dead end—an obstacle to clarity best avoided altogether—Bogoliubov adopted a more conservative position. As he put it in a 1955 paper:

> "[Perturbative QFT] best corresponds to the actual modern state of field theory, where up to the present various formal expansions in powers of the smallness of the interaction cannot be removed and where all fundamental results were obtained with the help of these expansions." [Bogoliubov, 1955].[14]

Bogoliubov's point was not that perturbation theory provided a complete or satisfactory description of quantum fields, but rather that it constituted the sole concrete source of knowledge about interacting QFTs available at the time. Given this situation, he argued, it was methodologically sensible to improve the mathematical clarity of the perturbative formalism, even if that formalism might ultimately prove inadequate. Improving rigour here was valued not as a guarantee of foundational security, but for its diagnostic potential: understanding how the theory worked in its current form could inform which of its assumptions might eventually require modification or abandonment.[15]

Although causal perturbation theory would later come to be understood as an ameliorative reconstruction of renormalised perturbation theory, its earliest incarnation was not conceived in this way. In the 1940s, its originator, Ernst Stueckelberg developed an alternative approach to relativistic quantum theory that occupied a middle position between Heisenberg's S-matrix programme and Dyson's perturbative framework. Like Heisenberg, Stueckelberg sought to avoid reliance on a local time-evolution equation, but like Dyson he also strove to rehabilitate the perturbative approximation scheme. His central idea was to derive a perturbative expansion for the S-matrix from a causality condition, rather than from the Schwinger–Tomonaga equation which Dyson had appealed to.

As in axiomatic QFT, this required a rethinking of the theory's basic objects. Stueckelberg retreated from Heisenberg's insistence on working solely with asymptotic quantities, but he did so not by returning to quantum fields, but by generalising his S-matrix operators. As noted in Section 2, Stueckelberg [1951] attempted to compute Dyson's finite-time evolution operator $U(t,t')$ and discovered a new class of divergences—so-called boundary divergences. These can be understood in distribution-theoretic terms: restricting the evolution operator to sharp time slices amounts to evaluating a singular distribution on

---

[14] Translation of this Russian source was provided by Kseniia Mohelsky.

[15] Indeed, Bogoliubov starts his paper with the following statement:

> In this situation many researchers seek to go beyond the frame of the problem by, for example, giving up rigorous locality or strict observance of the law of causality. That is why it is desirable to have a representation of quantum field theory that would allow us to see those basic physical assumptions, on which it is built in its recent form, in order to be able to understand, in which directions it is acceptable to generalize them. [Bogoliubov, 1955] (Kseniia Mohelsky's transation).

a hypersurface, which is mathematically ill-posed. Stueckelberg's solution was to replace $U(t,t')$ with *smeared* evolution operators $S(g)$, where the smooth function $g(x)$ plays the role of a switching function that localises the interaction in some space-time region. Formally, g functions as a test function, rendering the corresponding operator-valued distribution well defined.

This move parallels the first step of axiomatic QFT, where ill-defined pointlike fields are replaced by smeared operators. However, the two programmes diverge sharply in what they aim to do with these rehabilitated local objects. In axiomatic QFT, locality and causality conditions are used to derive structural results without constructing explicit model-dependent dynamics. In causal perturbation theory, by contrast, the aim was precisely to recover a concrete model-specific dynamics. Stueckelberg's causality condition was intended to replace the Schwinger–Tomonaga equation as the mechanism by which a perturbative expansion is generated.

It was Nicolay Bogoliubov who brought this idea to full fruition, recasting causal perturbation theory as a reconstruction of conventional renormalised perturbation theory rather than a rival framework. Using Stueckelberg's $S(g)$ operators, Bogoliubov formulated a precise causality condition. If two test functions g_1 and g_2 have supports in space-time regions G_1 and G_2, and no point of G_2 lies in the future light cone of any point of G_1, then the condition reads:

$$S(g_1 + g_2) = S(g_2)S(g_1). \tag{3}$$

Intuitively, this expresses the requirement that the operator associated with an earlier period of scattering acts on the state first; if the S(g) operator factorises, it factorises into a so-called time-ordered product. Importantly, this condition is stronger than the microcausality condition familiar from axiomatic QFT, which constrains operator products only at spacelike separation. Here, the causality condition plays a genuinely dynamical role, governing how operators compose in time without appeal to a local evolution equation.

Bogoliubov showed that this condition suffices to derive a perturbative expansion of the same general form as the Dyson series. One begins with an abstract expansion

$$S(g) = \sum_{n=0}^{\infty} \frac{1}{n!} \int S_n(x_1, \dots, x_n)\, g(x_1) \cdots g(x_n)\, \mathrm{d}x_1 \cdots \mathrm{d}x_n, \tag{4}$$

with the coefficients S_n initially unspecified. Imposing causality, together with unitarity and Poincaré invariance, yields recursive relations between the coefficients at successive orders. Fixing the first-order term,

$$S_1(x) = i\lambda H_I(x), \tag{5}$$

which amounts to selecting an interaction density, allows the higher-order terms to be constructed iteratively, reproducing the formal structure of the Dyson series.

This new derivation circumvented some of the problems with Dyson's original Schwinger-Tomonaga equation based derivation of the series. But perhaps more importantly, it set the stage for a more rigorous treatment of the ultraviolet divergences problem. It is at this stage that distribution theory explicitly enters the causal perturbation theory tradition—apparently independently of its adoption within axiomatic QFT.[16] Recall that, in the conventional formulation, divergent integrals appear from the second-order of the power series. Bogoliubov and Stueckelberg both reinterpreted these divergences as arising from a naïve multiplication of singular distributions. Unlike ordinary functions, distributions do not in general admit well-defined products, especially when their singular supports overlap—$\delta(x)^2$, for instance, is not well-defined. The $S_1(g(x))$ operators (or propagators in the conventional formulation) which are multiplied together in the higher-order terms of the expansion, are, in fact, singular distributions with overlapping support.

Since the causal derivation of the series basically treated the series coefficients as distributional from the beginning, it was uniquely positioned to articulate this diagnosis and present a formal solution. The ultraviolet divergences problem was thus recast in the following form: is it possible to construct the distributional products which appear in each term of the perturbation series? Bogoliubov and Parasiuk [1957] suggested a schema for carrying out this construction: first construct the product over a restricted space of test functions which vanish at the origin (essentially problematic area of overlapping singular support) and then extend this product to the full space of test functions. It turns out that this construction can be carried out—though it took some time for this to be carried out rigorously and explicitly in the later work of Epstein and Glaser [1973]. Thus, causal perturbation theory ultimately achieved something quite remarkable which remains little appreciated today. It showed that perturbative ultraviolet divergences do not in fact indicate the physical breakdown of QFT at high energies; rather, they are artefacts of a naive treatment of distributional products.

Let us close this section by situating these successes in the context of the programme's broader methodological strategy. Bogoliubov's reflections on what the programme had achieved in the mid-1960s are quite revealing:

> "We now know how to formulate a perturbation theory in which any fixed approximation is free from divergences. This certainly represents only a very *partial success*... [However] we do not have even the first intuitive idea of how to sum the whole series. Thus, the question of whether in general quantum electrodynamics even exists as a theory, free of internal contradictions, remains open." [Bogoliubov, 1965, 31–32] (my italics).

This notion of partial success, is characteristic of the local-repair rigourisation strategy that the programme enacted. Causal perturbation theory had de-fanged one of the most

---

[16] For a detailed discussion of how distribution theory entered the causal perturbation theory programme, see Blum and Fraser [2025].

notorious mysterious internal anomalies facing QFT—the ultraviolet divergences problem. However, it had left other internal problems entirely untouched. It did nothing to address the large-order divergence of the perturbation series, nor did it mitigate non-perturbative inconsistency worries such as the Landau pole.

Rather than attempting to secure global control over QFT in a single sweep, as axiomatic QFT had, the local approach works to isolate particular sources of mathematical unclarity and addresses them one-by-one, bracketing other foundational issues—and, indeed, bracketing the question of the framework's global consistency—in the process. This makes incremental progress possible, of a kind which is difficult to envisage in global-rebuilding approach. However, there is an accompanying risk of wasting time patching up a sinking ship.

The perception of perturbative QFT as a sinking ship was almost universal in the 1950s-1960s, likely contributing to the relative obscurity of causal perturbation theory, despite its apparent successes. As we shall discuss in the next section these evaluations of perturbation theory would shift drastically in the 1970s, though not in a way that straightforwardly vindicated the causal perturbation theory programme.

5 1970s: A New Problem Situation

In the preceding sections, I have presented axiomatic QFT and causal perturbation theory as alternative ways of mobilising rigour in response to the problem situation confronting quantum field theory in the 1950s. An important complication of this picture, however, is that the problem situation to which rigour-orientated reformulation programmes respond is itself historically dynamic. As the perceived risks and prospects associated with QFT shifted, so too did the methodological rationales underwriting these programmes.

The key turning point was the emergence, in the early 1970s, of quantum chromodynamics (QCD) as a successful theory of the strong interactions. This development involved a number of surprises that substantially altered the way QFT was viewed as a general theoretical framework. One was the discovery of asymptotic freedom. Whereas the Landau pole problem had fuelled fears that the renormalised couplings of interacting QFTs would diverge at high energies, QCD displayed precisely the opposite behaviour: its coupling strength decreases with increasing energy and tends to zero in the ultraviolet. This suggested that at least some of the most troubling inconsistency worries of the 1950s were not generic features of QFT as such, but model-specific pathologies. As a result, confidence in the global viability of the QFT framework increased markedly.

Equally striking was the role that perturbation theory came to play in QCD. Strong interactions had long been assumed to demand intrinsically non-perturbative methods. Yet asymptotic freedom made it possible to extract precise quantitative predictions from perturbative calculations at high energies. In this respect, the resurgence of perturbation theory was, if anything, more surprising than the rehabilitation of QFT itself. In retrospect, the axiomatic QFT tradition's insistence on non-perturbativity had been closely aligned with the prevailing phenomenological expectations of the 1950s and 1960s; by the 1970s, however, those expectations had shifted.

One might expect these developments to have created fertile ground for the local-repair strategy exemplified by causal perturbation theory. After all, if perturbation theory was now indispensable to nuclear phenomenology, then clarifying its mathematical foundations might seem more urgent than ever. While the changing landscape likely did sustain continued interest in the programme—most notably in the rigorous formulation of Epstein–Glaser renormalisation in the early 1970s—it did not trigger anything like a renaissance. Instead, activity in causal perturbation theory remained a minority pursuit, even within the small community of mathematical physics.

A plausible explanation for this is that the effective field theory (EFT) perspective, which rose to prominence alongside QCD, ate the problem space which causal perturbation theory had been reacting too. EFTs offered an epistemic rather than formal resolution of ultraviolet difficulties; if quantum field theories are understood as intrinsically low-energy descriptions, then ultraviolet divergences cease to be threats to consistency and instead become signals of the breakdown of a theory's domain of applicability. From this vantage point, a careful distribution-theoretic reconstruction of renormalisation no longer appeared methodologically pressing, even if it remained conceptually illuminating.

Axiomatic QFT, for its part, also underwent a process of realignment. One visible shift in the 1970s was an increased emphasis on the explicit construction of interacting models satisfying the axioms. Whereas model-independence had been a virtue during the earlier period of foundational uncertainty, a central ambition now became to construct realistic theories in four-dimensional Minkowski space-time (as evidenced by a Millennium Prize Problem which explicitly asks for a construction of a Yang–Mills theory in four dimensions [Jaffe and Witten, 2006]). While the earliest successes of what would become known as constructive field theory in fact go back to the late 1960s,[17] including the construction of non-trivial models in a reduced number of space-time dimension, the 1970s saw a shift in their perceived role, as the aim of constructing physically realistic models came to the fore. This goal has remained elusive, however, largely due to the severe technical difficulties posed by non-linear interactions in four dimensions.

At the same time, we can arguably identify a quieter rebranding of axiomatic QFT's broader methodological aspirations. As discussed in Section 3, the use of the axiomatic method in the 1950s and 1960s was motivated primarily by the need to secure conceptual and mathematical control in the face of paradox and inconsistency. In the post-QCD era, however, it became increasingly plausible to frame axiomatic QFT in more overtly Hilbertian terms—as a project aimed at isolating the minimal structural assumptions underlying a well-established branch of physics. Retrospective discussions like Wight-

---

[17] A historical puzzle therefore remains about why the constructive field theory movement originated when it did, some years before the discovery of asymptotic freedom.

man [1976] can be understood as developing this new self image rather than as neutral recollections. In this justificatory rather than crisis-driven register, axiomatic QFT could be presented as contributing to our understanding of why QFT works, rather than as a strategy for dealing with inconsistency threats.

It is possible, of course, to adopt a critical stance toward these evolving methodological rationalisations. One way of reading Wallace's critique of axiomatic QFT, for instance, is as an argument that the programme failed to articulate a compelling role for itself in the context of the effective field theory paradigm. My aim here, is not to assess the quality of underpinning methodological rationalisations, but to point to their adaptive character.

In line with this dynamical perspective, it is worth pointing out that the division between axiomatic QFT and causal perturbation theory did not persist indefinitely. From the 1990s onward, elements of the two traditions were explicitly recombined in what is now known as perturbative algebraic quantum field theory (see [Fraser and Rejzner, 2024]). This framework integrates the causal perturbation theory treatment of perturbative renormalisation with the algebraic structures characteristic of the Haag-Kastler axioms. While these recent developments deserve their own analysis, this can provisionally be understood as a historically delayed recombination made possible by a further transformed problem situation—one in which confidence in the global viability of QFT had stabilised, and in which perturbative methods had been re-legitimised by their central role in QCD and effective field theory. In retrospect, the independent developments of axiomatic QFT and causal perturbation theory may be less a matter of principle than of divergent assessments of the risks and opportunities each approach offers.

6 Global Rebuilding vs Local Repair

I conclude this chapter with some more general reflections on the methodological operationalisation of the rigour concept. Let us return to the question we started with: how should we account for the diversity of rigour-orientated reformulation programmes in QFT?

A naive approach to answering this question might be to view axiomatic QFT and causal perturbation theory as combining a commitment to improving rigour with different ancilliary commitments: specifically, a commitment to developing the perturbative formalism, in the case of causal perturbation theory, and a commitment to developing an entirely non-perturbative formalism, in the case of axiomatic QFT. This approach treats rigour as a univocal ideal which can be applied to different theoretical objects.

A more promising approach, in my view, is to instead view rigour as an ideal which is cashed out differently in the context of different methodological strategies. There will, in some cases, be different ways to make an area of mathematised science more rigorous, associated with different risk-opportunity profiles. This is the way that I suggest viewing the global-rebuilding and local-repair rigorousiation strategies which I have introduced to make sense of the development of axiomatic QFT and causal perturbation theory respectively.

An analogy which may help to conceptualise the opportunities and risks associated with these two methodological paths is the following. Imagine engineers confronted with a mysterious technological artefact whose observable performance is impressive but whose internal organisation is poorly understood, and which has begun to exhibit recurring malfunctions of unclear origin. One response—analogous to global-rebuilding—would be to attempt to reconstruct it from the ground up, producing a new version that replicates the salient functions of the original while resting on a transparent and fully controlled architecture. In certain respects, this strategy is safer: insofar as the reconstruction proceeds from well-specified components and design principles, each subsystem can be tested in isolation and the overall consistency of the new artifact secured by construction. Until the rebuilt device is completed and shown to function, however, little has been achieved in stabilising the behaviour of the original system (though arguably progress may still have been made in articulating the principles that such a system must satisfy).

Alternatively—by analogy with a local-repair strategy—engineers might intervene directly within the malfunctioning artifact, probing particular subsystems, diagnosing sources of failure, and introducing piecemeal fixes while leaving the overall architecture largely intact. This approach permits incremental advances: specific faults can be isolated, their causes understood, and limited improvements implemented without waiting for a wholesale redesign. Yet it also carries a different kind of risk. Time and effort may be expended refining components of a device whose deeper structural problems ultimately render it unviable, so that local successes fail to accumulate into a durable resolution.

The analogy can also be extended to illuminate the different kinds of understanding that the two strategies promise to deliver. Both routes may yield increased insight into the functioning of the artefact, but in importantly different senses. Only the local-repair strategy is well-suited to revealing what is wrong with this particular device: by isolating failing components and tracing the origins of specific malfunctions, it affords a form of diagnostic understanding that global-rebuilding cannot supply. This corresponds to the way in which causal perturbation theory clarifies the ultraviolet divergence problem by working inside the perturbative formalism in which that difficulty arises—something axiomatic reconstruction could not do precisely because it bracketed that apparatus from the outset.

By contrast, the process of global-rebuilding may yield a different kind of understanding, which one might dub architectural understanding. Reconstructing the artifact from new design principles can reveal which general structural features devices of this kind must possess in order to perform their characteristic functions at all. Moreover, rebuilding opens up the possibility not merely of reproducing the original unit, but also of improving on it, extending its capacities. This helps to explain how axiomatic QFT could, at certain moments in its history, sustain ambitions of underwriting genuinely non-perturbative techniques that would go beyond the reach of perturbation theory. Local-repair, by contrast, is not typically a route to such extensions in its own right (though it may indirectly inform future theorising by clarifying the limitations of the existing framework).

Of course, there is no reason in principle why tinkering with the existing unit and attempting to build a replacement in tandem might be fruitful, with each project informing the other. However, when both paths are judged to be extremely costly to progress, and the risks and payoffs of each remain deeply uncertain, one can see how the choice between them could be perceived as a sharp methodological fork. This, I suggest, is what happened with axiomatic QFT and causal perturbation theory in the 1950s, giving rise to the distinct traditions we see today.

Now that the global and local rigourisation strategies have been distinguished in more general terms, the question arises of whether this distinction can be usefully applied to other scientific contexts. I will conclude with some speculative comments in this direction.

One pure mathematics context which may be instructive to consider (and which also has a direct bearing on the case of QFT) is the theory of distributions. Schwartz [1951]'s book is rightly recognised as a foundational work in the theory of distributions. However, the soviet mathematician Sobolev [1936] had some decades earlier introduced generalised functions in the context of the analysis of partial differential equations (PDEs). Schwartz is sometimes characterised as taking this earlier work and making it rigorous. However, my discussion suggests an alternative narrative: Sobolev applied rigour locally, to treat specific problems that arose in the context of PDE's, while Schwartz applied rigour globally in order to develop a systematic structural theory of distributions. As Bogoliubov came from the same intellectual tradition as Sobolev, an intriguing possibility for historical investigation would be the extent to which different rigourisation strategies are favoured in different mathematical cultures.

What about philosophers' views of rigour? Philosophers of science have arguably tended to privilege global-rebuilding rigour in their analyses of scientific methodology. This is understandable given the formative role of logical empiricism on the discipline, which can be viewed as the epitome of the global-rebuilding schema. Global structure is also often understood to be implicated in questions about a theory's fundamental ontology, which have been central to the disciplinary identity of the philosophy of physics. The diagnostic insights brought about by local applications of rigour can undoubtedly be conceptually significant, however, as causal perturbation theory's resolution of the ultraviolet divergence problem attests to. One lesson from the preceding discussion, therefore, is that broadening our understanding of how rigour is enacted may philosophers to comprehend important scientific developments which would otherwise be obscured.

Acknowledgements

The ideas developed in this chapter emerged while working with Alexander Blum on the project that eventually became Blum and Fraser [2025]. I am very grateful to him for giving me his blessing to work them out here independently. Alexander Blum, Francisco Calderón and Porter Williams gave helpful comments on drafts of this work which informed the final version. I presented early version of this material, and received useful feedback, at the following events in 2025: "Mathematical Rigour in Physics" workshop, IHPST, Paris 1 Panthéon-Sorbonne, Relativistic Quantum Information COST online seminar, and 100 Years of Quantum: Perspectives on its Past, Present, and Future, Perimeter Institute for Theoretical Physics. I would also like to thank Kseniia Mohelsky for providing English translations of the Russian sources quoted from in this chapter. This research was supported by funding from the ASYMPTOPHYS project (ANR-22-CE54-0002, hosted by IHPST, CNRS, UMR8590, France).

References

Zoe Ashton. *A Social Epistemological Approach to Mathematical Rigor*. The Ohio State University, 2024.

Michael J Barany, Anne-Sandrine Paumier, and Jesper Lützen. From nancy to copenhagen to the world: The internationalization of laurent schwartz and his theory of distributions. *Historia Mathematica*, 44(4):367–394, 2017.

A. S. Blum. *Probing the Consistency of Quantum Field Theory I: From Nonconvergence to Haag's Theorem*. Cambridge University Press, 2026.

A. S. Blum. *Probing the Consistency of Quantum Field Theory II: From the Landau Pole to the Renormalization Group*. Cambridge University Press, Forthcoming.

A. S. Blum and J. D. Fraser. Perturbative causality. *Synthese*, 2025. Forthcoming.

N. Bogoliubov and O. Parasiuk. Über die Multiplikation der Kausalfunktionen in der Quantentheorie der Felder. *Acta Mathematica*, 97(1):227–266, 1957. ISSN 0001-5962.

N. Bogoliubov and D. Shirkov. *Introduction to the Theory of Quantized Fields*. Introduction to the Theory of Quantized Fields. Interscience, New York, 1959.

N. N. Bogoliubov. The causality condition in quantum field theory. *Izv. Akad. Nauk SSSR, Ser. Fiz.*, 19:237, 1955.

N. N. Bogoliubov. Mathematical problems of quantum field theory. *Russian Mathematical Surveys*, 20(3):28, 1965.

James T Cushing. *Theory construction and selection in modern physics: The S-matrix*. Cambridge University Press, 1990.

F. J. Dyson. Divergence of perturbation theory in quantum electrodynamics. *Physical Review*, 85(4):631, 1952.

Freeman J Dyson. The S-matrix in quantum electrodynamics. *Physical Review*, 75(11): 1736, 1949a.

Freeman J Dyson. The radiation theories of Tomonaga, Schwinger, and Feynman. *Physical Review*, 75(3):486, 1949b.

H. Epstein and V. Glaser. The role of locality in perturbation theory. *Annales de l'institut Henri Poincaré. Section A, Physique Théorique*, 19(3):211–295, 1973.

D. Fraser. How to take particle physics seriously: A further defence of axiomatic quantum field theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 42(2):126–135, 2011.

Doreen Fraser. Quantum field theory: Underdetermination, inconsistency, and idealization. *Philosophy of Science*, 76(4):536–567, 2009.

J. D. Fraser and K. Rejzner. Perturbative expansions and the foundations of quantum field theory. *The European Physical Journal H*, 49(1):10, 2024.

James D Fraser. The twin origins of renormalization group concepts. *Studies in History and Philosophy of Science Part A*, 89:114–128, 2021.

Marvin L Goldberger. Use of causality conditions in quantum theory. *Physical Review*, 97(2):508, 1955.

R. Haag and D. Kastler. An algebraic approach to quantum field theory. *Journal of Mathematical Physics*, 5(7):848–861, 1964.

Rudolf Haag. On quantum field theories. *Dan. Mat. Fys. Medd*, 29(12):1–37, 1955.

Rudolf Haag. Some people and some problems met in half a century of commitment to mathematical physics. *The European Physical Journal H*, 35(3):263–307, 2010.

W. Heisenberg. Die beobachtbaren grössen in der theorie der elementarteilchen. *Zeitschrift für Physik*, 120:513–538, 1943a.

W. Heisenberg. Die beobachtbaren grössen in der theorie der elementarteilchen II. *Zeitschrift für Physik*, 120:673–702, 1943b.

W. Heisenberg. Die beobachtbaren grössen in der theorie der elementarteilchen III. *Zeitschrift für Physik*, 123:93–112, 1944.

Werner Heisenberg and Wolfgang Pauli. Zur quantendynamik der wellenfelder. *Zeitschrift für Physik*, 56(1-2):1–61, 1929.

C. A. Hurst. The enumeration of graphs in the feynman-dyson technique. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 214(1116): 44–61, 1952.

Arthur Jaffe. Divergence of perturbation theory for bosons. *Communications in Mathematical Physics*, 1(2):127–149, 1965.

Arthur Jaffe and Edward Witten. Quantum yang-mills theory. *The millennium prize problems*, 1:129, 2006.

Gunnar Källén. *On the magnitude of the renormalization constants in quantum electrodynamics*. Munksgaard, 1953.

L. Klein. *Dispersion relations and the abstract approach to field theory*. Gordon and Breach Publishers, New York, 1961.

Lev D Landau, A Abrikosov, and L Halatnikov. On the quantum theory of fields. *Il Nuovo Cimento (1955-1965)*, 3:80–104, 1956.

Lev Davidovich Landau and II Pomeranchuk. On point interactions in quantum electrodynamics. In *Dokl. Akad. Nauk SSSR*, volume 102, page 489, 1955.

Harry Lehmann, Kurt Symanzik, and Wolfhart Zimmermann. Zur formulierung quantisierter feldtheorien. *Il Nuovo Cimento (1955-1965)*, 1(1):205–225, 1955.

Wolfgang Pauli. The connection between spin and statistics. *Physical Review*, 58(8):716, 1940.

Karl R. Popper. *Objective Knowledge: An Evolutionary Approach*. Clarendon Press, Oxford University Press, Oxford, United Kingdom, 1972.

Miklós Rédei. Hilbert's 6th problem and axiomatic quantum field theory. *Perspectives on Science*, 22(1):80–97, 2014.

Miklós Rédei and Michael Stöltzner. Soft axiomatisation: John von neumann on method and von neumann's method in the physical sciences. In *Intuition and the axiomatic method*, pages 235–249. Springer, 2006.

Pablo Ruiz de Olano, James D Fraser, Rocco Gaudenzi, and Alexander S Blum. Taking approximations seriously: The cases of the chew and nambu-jona-lasinio models. *Studies in History and Philosophy of Science*, 93:82–95, 2022.

Laurent Schwartz. *Théorie des distributions*. Hermann, Paris, 1951. 2 volumes.

Irving E Segal. The mathematical meaning of operationalism in quantum mechanics. In *Studies in Logic and the Foundations of Mathematics*, volume 27, pages 341–352. Elsevier, 1959.

Suman Seth. Crisis and the construction of modern theoretical physics. *The British Journal for the History of Science*, 40(1):25–51, 2007.

Sergei Sobolev. Méthode nouvelle à résoudre le problème de cauchy pour les équations linéaires hyperboliques normales. *Mar. Sbornik (N.S.)*, 1(1):39–72, 1936.

Michael Stöltzner. Opportunistic axiomatics—von neumann on the methodology of mathematical physics. In *John von Neumann and the foundations of quantum physics*, pages 35–62. Springer, 2001.

Michael Stöltzner. On optimism and opportunism in applied mathematics. *Erkenntnis*, 60(1):121–145, 2004.

Raymond F. Streater and Arthur S. Wightman. *PCT, Spin and Statistics, and All That*, volume 3 of *Mathematical Physics Monograph Series.* W. A. Benjamin, New York, 1964.

E. Stueckelberg and A. Petermann. La normalisation des constantes dans la théorie des quanta. *Helv. Phys. Acta*, 26:499–520, 1953.

E. C. G. Stueckelberg. Relativistic quantum theory for finite time intervals. *Phys. Rev.*, 81:130–133, 1951.

Fenner Stanley Tanswell. *Mathematical rigour and informal proof.* Cambridge University Press, 2024.

W. Thirring. On the divergence of perturbation theory for quantized fields. *Helvetica Physica Acta*, 26:33, 1953.

D. Wallace. Taking particle physics seriously: A critique of the algebraic approach to quantum field theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 42(2):116–125, 2011.

David Wallace. In defence of naiveté: The conceptual status of lagrangian quantum field theory. *Synthese*, 151(1):33–80, 2006.

A. S. Wightman. Hilbert's sixth problem: Mathematical treatment of the axioms of physics. In F. E. Browder, editor, *Mathematical Developments Arising from Hilbert Problems*, volume 28 of *Symposia in Pure Mathematics*, pages 147–240. 1976.

Arthur S. Wightman. Interview with jagdish mehra. *Unpublished interview transcript. Interview conducted 24 Feb 1987; transcript dated 18 Jan 1988*, 1987.

Arthur S Wightman. How it was learned that quantized fields are operator-valued distributions. *Fortschritte der Physik/Progress of Physics*, 44(2):143–178, 1996.