Synthese (2024) 204:116
https://doi.org/10.1007/s11229-024-04750-5

ORIGINAL RESEARCH

Staying on-shell: manifest properties and reformulations in particle physics

Josh Hunt[1]

Received: 31 December 2023 / Accepted: 16 August 2024

Abstract
The empirical success of particle physics rests largely on an approximation method: perturbation theory. Yet even within perturbative quantum field theory, there are a variety of different formulations. This variety teaches us that reformulating approximation methods can provide a tremendous source of progress in science. Along with enabling the solution of otherwise intractable problems, reformulations clarify what we need to know to obtain solutions, which can in turn make previously hidden properties manifest. To develop these lessons, I compare and contrast three compatible formulations of perturbative QFT: (i) elementary perturbation theory, (ii) the method of Feynman diagrams, and (iii) a recent reformulation known as on-shell recursion. I propose and defend a novel account of what it means to 'make a property manifest,' based on the inferences that a formulation warrants.

Keywords Reformulation · Approximation · Intellectual significance · Manifest properties · Feynman diagrams · Scattering amplitudes

1 Introduction

Many results in quantum field theory arise from an application of perturbation theory. Not just in QFT, but across physics and chemistry, perturbation theory underwrites important approximation methods—sometimes, the most important. Far from being an undifferentiated monolith, perturbation theory comprises a multifaceted collection of problem-solving techniques, unified by an overarching problem-solving strategy. In many cases, how we *formulate* perturbation theory makes a huge difference, both conceptually and methodologically.[1]

While philosophers of science commonly talk about different formulations of scientific theories, it is less common to focus on formulations of approximation methods.

[1] For discussions of the importance of perturbation theory for physics, see Batterman (2002), Fillion (2021), Ruiz de Olano et al. (2022, p. 85), and Miller (2023).

✉ Josh Hunt
jhunt11@syr.edu

1 Syracuse University, 541 Hall of Languages, Syracuse, NY 13210, USA

Published online: 19 September 2024

This is what I propose to do in this essay. I will focus on three different formulations of perturbative methods for quantum field theory: (i) elementary perturbation theory, (ii) Feynman diagrams, and (iii) on-shell recursion. At each stage of reformulating, we will see how new concepts—both mathematical and physical—clarify what we need to know to solve problems. This case study thereby supports and extends my 'conceptualist' account of reformulating (Hunt, 2021b, 2023). Additionally, this case study supports Ruiz de Olano et al.'s claim that making sense of high-energy physics requires "acknowledging approximation construction as a distinct stream of theoretical development," namely, an activity distinct from constructing new models or theories (2022, p. 88).

Section 2 introduces key notions from my account of reformulations, including compatible formulations, problem-solving plans, epistemic dependence relations, and methodological vs. intellectual significance. Section 3 establishes a baseline for comparison by introducing an elementary formulation of perturbation theory. I discuss the well-known Feynman diagram formulation in Sect. 4. Limitations of this approach motivate a more recent formulation known as *on-shell recursion*, which I introduce in Sect. 5. Section 6 expounds some epistemic benefits that on-shell recursion provides. This includes making manifest some properties that the Feynman diagram formulation obscures. In Sect. 7, I propose and apply a novel philosophical account of what it means to make a property manifest. On this account, a formulation makes a property manifest whenever it warrants inferring that the property obtains. I then proceed to gradate this account: relative to a given evidence set and problem of interest, one formulation makes a solution *more manifest* provided it rules out more epistemically possible but non-actual solutions than an alternative formulation does.

2 Compatible reformulations of approximation methods

In previous work, I have proposed and defended *conceptualism* as an account of how scientific reformulations improve understanding (2021a, 2021b, 2023). According to conceptualism, reformulations improve understanding by clarifying what we need to know to solve problems. Although this prior work does not explicitly highlight reformulating approximation methods, many of my examples involve just this, including reformulations of matrix element calculations in atomic physics (2021a) and reformulations of crystal field theory in quantum chemistry (2023). Indeed, like the present case study, both of these examples involve reformulating perturbation theory.

In conjunction with a theory or model, an approximation method provides a problem-solving plan. A plan consists of a series of steps, connected by inference rules. Each step of a problem-solving plan relates input information to output information, via an inference rule. For convenience, I refer to these components of plans as *epistemic dependence relations* or 'EDRs' for short (2023). EDRs specify what we need to know or what suffices to know to carry out a step in a problem-solving plan, serving as basic components in the epistemic structure of problem-solving. Significant reformulations involve changes to EDRs: in providing a new problem-solving plan, a non-trivial reformulation requires a different set of things we need to know (or that suffice to know) to solve a problem. In Lagrangian mechanics for instance, we learn that for many systems, we do not need to know the constraint forces in order to solve for the equations of motion. In this way, Lagrangian and Newtonian mechanics provide problem-solving plans based on different epistemic dependence relations.

Here, I focus on different calculational methods for approximating scattering amplitudes. Since these methods constitute logically consistent plans for solving a given problem—or, more generally, a class of problems—they function as *compatible formulations*. Compatible formulations do not disagree about the way the world is. They provide the same answers to problems within their shared domain of applicability. For such problems, they make the same predictions and ascribe compatible content to physical systems.[2] Throughout, I will use 'formulation' as a convenient shorthand for the more precise—but cumbersome—'problem-solving plan.' Differences in EDRs arise naturally in the context of formulations of perturbation theory. Using on-shell recursion, we learn that for many quantum field theories, we do not need to know the Feynman diagrams corresponding to a given tree-level amplitude. Instead, as Section 5 illustrates, it suffices to know a finite number of 'seed amplitudes' along with recursion relations. In such cases, "the mere knowledge of three-point [seed] amplitudes allows the construction of *all* higher point amplitudes in a recursive fashion" (Henn & Plefka, 2014, p. 36).

As noted above, a reformulation involves a change in problem-solving plan, sometimes leading to an entirely novel approach. While some reformulations are worthwhile, others add little to nothing of value, and still others make matters worse. In the case study that follows, I identify various ways that a reformulation can constitute scientific progress. In particular, I focus on two dimensions of progress: the 'methodological' and the 'intellectual.' On the methodological side, good reformulations often lead to calculations that are shorter, easier, or reduce the risk of mistakes. On the intellectual side, good reformulations isolate independently-treatable aspects of a problem, unify different phenomena, uniformly treat different problems, or illuminate patterns that are otherwise hidden or surprising. Of course, to call a reformulation 'good' (or 'better') is to make a judgment of its (comparative) value. Such judgments seemingly commit us to facts about the value of reformulations. Here, I remain neutral on how best to understand these normative or evaluative commitments.[3]

[2] Ruiz de Olano et al. (2022, pp. 86–88) provide an extended illustration of how some uses of perturbation theory can lead to *incompatible* approximation methods, i.e. methods that make different empirical predictions or ascribe logically-incompatible content to a physical system. As Ruiz de Olano et al. (2022) note, such cases seem to lead to genuinely *different* or new approximation methods, rather than what I call reformulations of the *same* approximation method. I set aside here further questions about how best to individuate approximation methods.

[3] Regarding my own meta-normative commitments, I endorse expressivism about the goodness or value of reformulations. When we judge that a reformulation is either (intellectually) good or better than another, we at least express acceptance of a set of norms that recommend (intellectually) preferring that formulation.

3 Elementary perturbation theory

An n-point scattering amplitude 'A_n' characterizes the likelihood that a given particle interaction will take place, where n-many particles are involved.[4] Amplitudes provide a bridge between theoretical predictions and empirical observation. The norm-squared of the amplitude ($|A_n|^2$) is proportional to the differential cross-section, which quantifies how likely a given scattering process is to occur. Owing to the central importance of scattering amplitudes for characterizing the empirical content of a QFT, physicists devote a great deal of theoretical resources to calculating them for various particle interactions.

Formally, an amplitude A_n is an inner product of an initial particle state '$|in\rangle$' and a final particle state '$|out\rangle$'. Due to computational complexities, we typically approximate scattering amplitudes by applying perturbation theory. To calculate a scattering amplitude, we treat it as a matrix element of a scattering matrix S. This matrix relates idealized in- and out-states at negative and positive temporal infinity, respectively. In these limits, we can idealize the relevant particles as non-interacting, enabling us to use a non-interacting Hamiltonian H_0. We then perturb this Hamiltonian by an interaction term H' to model particle scattering.[5]

To approximate a scattering matrix, we can expand it in powers of the coupling between quantum fields, relying on a formula known as Dyson's expansion.[6] We can then compute the amplitude to a desired order in the coupling by summing the terms that contribute at this order. Calculating these individual terms—known as time-ordered vacuum expectation values ('VEVs')—can itself be complicated. One method is to apply Wick's theorem, which re-expresses time-ordered VEVs in terms of quantities that are easier to calculate, namely normal-ordered operators and Wick contractions of pairs of operators.[7] Wick developed this theorem in Wick (1950), building on the special case of contractions between bosonic fields introduced by Houriet and Kind (1949).[8]

[4] More precisely, a scattering amplitude is a complex-valued function that acts on the "external data" which defines a scattering process (comprising the energies and momenta of the initial and final particles).

[5] For a brief overview of this approach, see Fraser (2020, p. 395ff.). Other formulations exist for carrying out elementary perturbation theory, such as expressing a generating functional as a path integral. See Srednicki (2007).

[6] Dyson's expansion expresses the time-evolution operator $U(t_2, t_1)$ (in the interaction picture) as a time-ordered exponential of the interaction Hamiltonian H' integrated over time: $U(t_2, t_1) = $ Time-order $[e^{-i\int_{t_1}^{t_2} H'(t)dt}]$. We then relate the scattering matrix to U in the infinite time limit (i.e. t_2, t_1 approach ∞ and $-\infty$, respectively) and approximate the exponential by expanding in powers of the interaction term H'. See Lancaster and Blundell (2014, §18) for details.

[7] Whereas time-ordering rearranges the operators in a vacuum expectation value so that those defined at later times precede those defined at earlier times, normal-ordering rearranges them such that all creation operators precede all annihilation operators. A Wick contraction of two operators is the difference between their time-ordered and normal-ordered values. See Rivat (2021, p. 12129) for standard interpretations of VEVs.

[8] Thanks to Alexander Blum for suggesting this paper by Houriet and Kind. My discussion deviates from the historical progression of these reformulations, since Feynman introduced his famous diagrams in 1948, *before* the introduction of Wick contractions. This element of rational reconstruction does not threaten my philosophical theses. For the historical development of Feynman diagrams, see Kaiser (2005) and Wüthrich (2010).

Already within this elementary approach to perturbation theory, an opportunity for a simple reformulation arises. Naïvely, one must consider *all possible* Wick contractions when calculating VEVs of time-ordered operators. However, one can show that any VEV term that has non-contracted operators equals zero—i.e. whenever there is at least one operator in a VEV not paired with any other by a Wick contraction, that VEV vanishes. This fact leads to the following epistemic dependence relation (EDR): in order to compute an amplitude (at a given order in perturbation theory), it suffices to calculate Wick-contracted terms where all operators are contracted (i.e. paired off). Hence, all VEVs with an *odd* number of operators necessarily vanish, since it is impossible to pair off an odd number of operators. This EDR focuses attention on a much smaller number of terms, illustrating a common methodological benefit of reformulating: by learning in advance that some terms of a calculation necessarily vanish, we free ourselves from the labor of calculating them.

Nevertheless, as described here, elementary perturbation theory requires us to calculate many terms that contribute only to the trivial part of the scattering amplitude S. This trivial part characterizes non-interacting particle processes, but our interest is in the non-trivial part characterizing interactions between particles.[9] Both epistemically and methodologically, we would prefer a formulation that focuses attention on the terms that contribute non-trivially. This would not only save us calculational labor, but also clarify what we need to know to describe interaction processes. The method of Feynman diagrams provides one reformulation of perturbative QFT that accomplishes these goals.

4 Feynman diagrams

Using Feynman diagrams, we can reformulate perturbation theory to focus attention on only those terms that contribute non-trivially to the scattering amplitude. This approach establishes a correspondence between diagrams and terms in the perturbative expansion. Each Feynman diagram is either path-connected or disconnected, depending on whether one can reach any given point in the diagram from any other point following a path of propagator lines. This topological property allows us to identify which terms contribute non-trivially. In short, Feynman diagrams allow us to take advantage of the following epistemic dependence relation: to approximate the scattering matrix, it suffices to calculate *connected terms*, where a connected term is one corresponding to a path-connected Feynman diagram (Srednicki, 2007, p. 65). Feynman diagrams are not the only way to express this EDR, but they are the most commonly used expressive means for taking advantage of it.[10]

[9] The scattering matrix S decomposes into a trivial part (the identity operator '1') characterizing non-interactions and a non-trivial part (T) characterizing interactions between particles: $S = 1 + iT$.

[10] Both (i) Wick contractions of vacuum expectation values and (ii) position space representations of terms in the perturbative expansion have the resources to express whether a term is connected or disconnected. Yet for a variety of reasons, they are at least practically less convenient for agents like us to structure the search space of identifying all path-connected terms at a given order of perturbation theory. I borrow the notion of 'expressive means' from Ken Manders.

The Feynman diagram reformulation illustrates some of the intellectual benefits that conceptualism is designed to capture. By taking advantage of topological properties of the terms, we learn that many terms need not be calculated. In other words, we clarify what we need to know to approximate the scattering matrix at a given order in perturbation theory. As noted above, to calculate the non-trivial part of S, one only needs to calculate the terms corresponding to connected diagrams. Moreover, a subset of connected diagrams—known as 'tadpoles'—can also be neglected. In this way, Feynman diagrams make manifest particular patterns already present in the perturbation series, picking out those terms that contribute to the non-trivial part of the scattering amplitude. They thereby provide a reformulation that is not only methodologically advantageous but also intellectually significant. Section 7 returns to this point by developing an account of what it means to make a property manifest.

4.1 Limitations of Feynman diagrams

In principle, we could calculate any scattering amplitude using Feynman diagrams (at least for theories described by a Lagrangian). But in practice, the Feynman diagram approach becomes incredibly complicated even for relatively simple scattering processes. A key reason is that the number of Feynman diagrams grows factorially even as the number of particles involved grows linearly (Henn & Plefka, 2014, p. vii). For gluon scattering at tree-level (i.e. first-order in perturbation theory), four-particle scattering involves four diagrams, five-particle scattering involves 25 diagrams, six particles require 220 diagrams, and ten particles require more than 1 million diagrams (Elvang & Huang, 2015, p. 8).[11] Hence, using Feynman diagrams for processes like these quickly becomes impractical.

Moreover, the desire to calculate such higher-point amplitudes is not purely theoretical; it matters for experiments. At the Large Hadron Collider, interaction processes are dominated by quarks and gluons scattering off each other and other particles, such as Z and $W^{\pm}$ bosons. These processes are known as *multiple jet events*, described by quantum chromodynamics (QCD). These generic processes form the backgrounds of most scattering signals. Hence, to isolate more interesting signals, these background processes must be calculated precisely so that they can be eliminated as noise. Detecting new, interesting signals requires calculating these backgrounds to third-order in perturbation theory, which includes content associated with two-loop Feynman diagrams (Cordero et al., 2022, p. 4; Dixon, 2016).[12] Future particle colliders may require calculating some backgrounds to fourth and fifth order in perturbation theory (Arkani-Hamed et al., 2021, p. 66). This requires calculating a multitude of terms at various orders in perturbation theory and provides practical motivation for developing on-shell methods, which are more tractable or computationally efficient than many alternatives. Historically, these computational motivations led first to *off-shell recursion* (Berends

[11] Badger (2016, p. 2) compares the rates at which the number of diagrams grows based on Feynman diagrams, color-ordered diagrams, and on-shell diagrams from BCFW recursion.

[12] For additional reviews of the current state of high precision calculations, see Heinrich (2021) and Campbell and Ellis (2023). Abreu et al. (2019) provides a representative example.

& Giele, 1988), which has many of the same methodological and epistemic benefits as on-shell recursion (introduced in Sect. 5).[13]

One might wonder: why can't we simply stick with Feynman diagrams but compute them numerically using a computer program, thereby sufficiently speeding up calculations to achieve our practical goals? Indeed, many of the most precise perturbative calculations have been done in exactly this way, using supercomputers.[14] However, since terms have to be properly canceled between diagrams, such calculations can be sensitive to the way in which the program handles poles. Even a computer calculation can go wrong.[15] Thus, working with thousands of diagrams is not only methodologically inconvenient but also epistemically risky. Here we see an important *practical* dimension of the epistemic value of a formulation: a compatible reformulation can reduce the likelihood of mistakes. The Feynman diagram calculations are error-prone in a way that on-shell recursion is not. Hence, on-shell recursion reduces the likelihood of making a mistake for both non-ideal agents and non-ideal computers. Nevertheless, in principle, Feynman diagrams contain the relevant information for making a successful calculation, even if they become impractical or risky.

Motivated by practical necessity, physicists in the 1980s developed strategies for making Feynman diagram calculations more tractable. These include the methods of helicity amplitudes (Gastmans & Wu, 1990) and color-ordered amplitudes. Both methods feature in the more sophisticated on-shell reformulation, but on their own, the number of Feynman diagrams still grows too quickly for feasible computer computations, especially at higher orders in perturbation theory. Section 5 returns to helicity amplitudes after introducing spinor–helicity variables. To color-order an amplitude, we factor out all $SU(N)$ gauge group degrees of freedom, resulting in *color-ordered* (a.k.a. *partial*) *amplitudes*. The full amplitude consists of partial amplitudes adjoined with the color degrees of freedom. For instance, in pure Yang–Mills theory at first-order in perturbation theory, all color degrees of freedom can be expressed as a sum over traces of the generator matrices T^a for the gauge group, leading to the following relationship between the full tree-level amplitude for n-many gluons, $\mathcal{A}^{tree}_{\text{n-gluon}}(\{a_i, p_i, h_i\})$, and the partial amplitudes $A^{tree}_{g^n}(\{p_i, h_i\})$. The latter are a function only of the kinematical information involving particle momenta and polarization states (note that the coupling

[13] For instance, Berends and Giele note that "the advantage" of recursion "is that for the calculation of an $n+1$ gluon process one can use the calculation of the n-gluon process.…the recursion relation takes automatically into account all Feynman diagrams. Writing down those diagrams would be a problem in itself, which is now avoided" (1988, p. 760). Badger et al. (2013a) compare the efficiency of on-shell vs. off-shell recursion.

[14] See, for instance, the numerical calculation of the electron's anomalous magnetic moment up to fifth-order in perturbation theory, involving thousands of Feynman diagrams (Aoyama et al., 2012). For a schematic introduction to this research program, see Kinoshita (1989). Miller (2023, p. 515) provides a brief overview. Note that numerical calculation introduces another approximation method beyond perturbation theory, such as Monte Carlo integration.

[15] The following remarks summarize some risks involved: "using computers to do the calculation can of course be very helpful, but not in all cases. Sometimes numerical evaluation of Feynman diagrams is simply so slow that it is not realistic to do. Moreover, given that there are poles that can cancel between diagrams, big numerical errors can arise in such evaluations. Therefore compact analytic expressions for the amplitudes are very useful in practical applications" (Elvang & Huang, 2015, p. 8). See also Abreu et al. (2019), Heinrich (2021) and Huss et al. (2023).

constant 'g' is removed from the partial amplitudes):[16]

$$\mathcal{A}^{tree}_{\text{n-gluon}}(\{a_i, p_i, h_i\}) = g^{n-2} \sum_{S_{n-1}} Tr(T^{a_{\sigma_1}} T^{a_{\sigma_2}} \dots T^{a_{\sigma_n}}) A^{tree}_{g^n}(\{p_i, h_i\}) \quad (1)$$

Color-ordering provides another illustration of how reformulating can be intellectually significant. Specifically, color-ordering *modularizes* an amplitude calculation into separate calculations for each color-ordered, partial amplitude. Using the appropriate symmetry factors, we can then combine these partial amplitudes into the full amplitude. Elsewhere, I have characterized 'modularization' as the decomposition of a problem into independently treatable sub-problems (2021a; 2023). Here, we see that color-ordering parallels the Wigner–Eckart theorem in atomic and molecular physics (2021a). By taking advantage of a system's symmetry, this theorem modularizes matrix element calculations into a symmetry-related component (Clebsch–Gordan coefficients) and a symmetry-invariant component (known as the reduced matrix element). This modularization teaches us that many matrix elements are related to each other by symmetries. A similar moral applies to color-ordering: to calculate the full amplitude, it suffices to calculate the color-ordered, partial amplitude and then later apply knowledge that depends only on symmetry properties of the amplitude. This separation of degrees of freedom improves our understanding of scattering amplitudes. On-shell recursion also takes advantage of color-ordering, making this technique applicable across different formulations of perturbation theory.

5 On-shell recursion

The calculational limitations of Feynman diagrams helped motivate physicists to reformulate the perturbative approach to QFT. Using a procedure known as *on-shell recursion*, physicists construct recursion relations that express a given amplitude in terms of amplitudes involving fewer particles. This reformulation involves multiple steps, including (i) re-expressing amplitudes in terms of spinor–helicity variables, (ii) applying Lorentz invariance, dimensional analysis, and a locality principle to determine the initial 'seed amplitudes' for recursion, and (iii) applying complex analysis and a principle of unitarity to derive recursion relations. Recursion provides a particularly powerful epistemic dependence relation: knowledge of the lower-point amplitudes yields knowledge of amplitudes involving arbitrarily-many particles. Because this formulation remains under-explored in the philosophical literature, this section provides a brief overview. More details can be found in the appendix. For simplicity, I focus on the case of scattering massless gluons ('pure Yang–Mills' theory), which conceptually forms the basis for scattering particles of arbitrary mass and spin (Arkani-Hamed et al., 2021; Liu et al., 2023).[17]

[16] See Henn and Plefka (2014, p. 23). Here, S_{n-1} is the permutation group on $n-1$-many objects. For more on color-ordering, see Schuster (2014) or Elvang and Huang (2015, p. 30ff.).

[17] For pedagogical introductions to the case of massless particles, see Elvang and Huang (2015) or Henn and Plefka (2014). For an introduction to the case of massive particles, see Ochirov (2018). The appendix discusses the scope of on-shell methods in more detail. Cushing mentions how in the early 1960s, Gunson

On-shell recursion and the Feynman diagram approach provide compatible formulations of perturbation theory. In the context of a specific quantum field theory, these two approaches agree on the models' symmetries, particle content, and scattering amplitudes (and hence all physical observables). Alongside elementary perturbation theory, I take them to be both predictively and ontologically compatible, rather than rival interpretations of quantum field theory or rival formulations of perturbation theory. For instance, physicists working in the on-shell framework frequently move back and forth between reasoning in terms of Feynman diagrams and reasoning in terms of on-shell expressions, indicating compatibility. Overall then, these three formulations provide compatible plans for approximating the scattering amplitudes of a quantum field theory.[18]

As the name suggests, on-shell recursion assumes that all particles are *on-shell*, i.e. on the "mass shell." This means that they satisfy a key kinematic constraint from special relativity, the energy–momentum relation $E^2 = c^4 m^2 + c^2 |p|^2$. Expressed in 4-vector notation in natural units: $p^\mu p_\mu = m^2$.[19] For instance, massless on-shell particles lie on the light cone. In contrast, off-shell particles—such as virtual particles represented by internal lines in a Feynman diagram—do not satisfy this relativistic constraint. They are in that sense unphysical (and hence typically not construed as literally representing particles).

The on-shell reformulation begins with a variable change to spinor–helicity variables, a "highly convenient and powerful notational tool" for expressing the amplitudes of massless particles for theories set in four spacetime dimensions and some other dimensions as well (Elvang & Huang, 2015, p. 15). According to Cheung, "spinor helicity variables are nothing more than an algebraic reshuffling of the external kinematic data. Such a manipulation would not be particularly advantageous were it not for the fact that scattering amplitudes enjoy an immense reduction in complexity when translated into these variables" (2017, p. 8). When expressed using these variables, many matrix elements of the scattering matrix vanish. These variables are thus akin to a good choice of basis, providing similar benefits to diagonalizing a matrix by choosing good basis functions in linear algebra. Additionally, these variables provide an expressive means that is both on-shell and Lorentz invariant, along with being gauge invariant.[20]

For a given particle, helicity is the projection of its spin along momentum: $\hat{h} = (\sigma \cdot \hat{\mathbf{p}}) / |\mathbf{p}|$, where $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \sigma_3)$ is the usual set of Pauli spin matrices. For a positive helicity particle, spin and momentum are aligned, whereas for a negative helicity particle they are anti-aligned. Although helicity is neither Lorentz nor gauge

(1965) gave an on-shell formulation of S-matrix theory, which Gell-Mann had suggested by at least 1956 (1990, pp. 185, 117, 81). Rivat (2023, §6) briefly discusses the on-shell formalism.

18 Different physical interpretations of these formulations could yield different verdicts about their ontological compatibility. If one interprets Feynman diagrams—but not on-shell recursion—as committed to virtual particles, then these formulations would no longer qualify as fully compatible. Elsewhere, I argue that on-shell recursion provides evidence that virtual particles are a formulational artifact of Feynman diagrams and hence should not be construed literally.

19 Throughout, I will use the mostly-minus metric convention: $\eta_{\mu\nu} = \text{diag}\,(+1, -1, -1, -1)$.

20 For an early application of spinor–helicity variables to the calculation of gluon scattering amplitudes in QCD, see Berends and Giele (1987). For a quick review of this formalism, in a variety of notations, see Elvang and Huang (2015, pp. 16-18) or Ochirov (2018, p. 2ff.).

invariant for massive particles, for massless particles it is, since their helicity coincides with their chirality.[21] Moreover, at high energies, many particles are approximately massless, including light quarks, charged leptons, and neutrinos (Gastmans & Wu, 1990, p. 3).[22] Hence, at high-enough energies, particle helicity is approximately conserved. Working in the helicity basis thereby simplifies matrix element calculations: matrix elements associated with processes that do not conserve helicity are guaranteed to vanish, providing a helicity-based selection rule.

To arrive at spinor–helicity variables, we begin with the usual momentum 4-vector p^μ. We re-express these four degrees of freedom as a momentum matrix $p_{a\dot{b}}$:

$$p_{a\dot{b}} = p_\mu \sigma^\mu_{a\dot{b}} = \begin{pmatrix} p_0 + p_3 & p_1 - ip_2 \\ p_1 + ip_2 & p_0 - p_3 \end{pmatrix} \tag{2}$$

Here, $\sigma^\mu = (1, \boldsymbol{\sigma})$. Notice that the determinant of this momentum matrix is $p^\mu p_\mu$, which for a massless particle equals zero. Hence, for massless on-shell particles, $det(p_{a\dot{b}}) = 0$, which implies that this 2×2 matrix has rank one. Applying an elementary fact from linear algebra, we can therefore express $p_{a\dot{b}}$ as an outer product of two 2-component vectors. These vectors constitute our *spinor–helicity* variables, denoted using square and angle bra-ket notation:

$$p_{a\dot{b}} = |p]_a \, \langle p|_{\dot{b}} \text{ and } p^{\dot{a}b} = |p\rangle^{\dot{a}} \, [p|^b \tag{3}$$

Spinor–helicity variables encode the helicity of spinor states: angle brackets correspond to negative helicity spinors, whereas square brackets correspond to positive helicity spinors.

We can express any function of kinematic data using inner products of angle and square spinors. For instance, the familiar Mandelstam invariants $s_{ij} = (p_i + p_j)^2$ from the Feynman diagram approach are easily re-expressed as follows:

$$s_{ij} = (p_i + p_j)^2 = 2p_i p_j = \langle p_i p_j \rangle \left[p_i p_j \right] \tag{4}$$

In general, by taking inner products of these angle and square spinor–helicity variables, we form Lorentz invariant building blocks: $\langle pq \rangle = \langle p|_{\dot{a}} \, |q\rangle^{\dot{a}}$ and $[pq] = [p|_a \, |q]^a$.

Rather than approximate the non-Lorentz invariant Feynman amplitude $A^{\mu_1 \cdots \mu_n}(p_1, \ldots, p_n)$ as a sum of Feynman diagrams, on-shell recursion aims to directly cal-

[21] For massive particles, a Lorentz boost can reverse momentum without reversing spin, making chirality the relevant Lorentz invariant for massive particles. In contrast, since massless particles move at the speed of light, it is impossible to change the direction of their momentum using a Lorentz boost. Hence, the helicity of a massless particle coincides with its chirality, which is invariant under gauge interactions. Right-handed massless particles have positive helicity, and left-handed particles have negative helicity (the opposite is true of anti-particles).

[22] Even at a time when particle colliders were much lower energy than today, Gastmans and Wu noted that at high energies, the masses of leptons and quarks (excluding perhaps the top quark) “are so small compared to the energies involved in the collisions that they can safely be neglected in almost all cases” (Gastmans & Wu, 1990, p. 3). Whether or not to count this massless limit as an abstraction, idealization, or approximation depends on one’s background philosophical views. See Fletcher (2019) and Morrison (2015, pp. 20–21).

culate *helicity amplitudes* $A_n(1^{h_1} \cdots n^{h_n})$.[23] Helicity amplitudes are Lorentz-invariant and gauge-invariant quantities, related to a product of polarization vectors ϵ^{h_i} with a corresponding Feynman amplitude: $A_n(1^{h_1} \cdots n^{h_n}) = \epsilon^{h_1}_{\mu_1} \ldots \epsilon^{h_n}_{\mu_n} A^{\mu_1 \cdots \mu_n}(p_1, \ldots, p_n)$. Each helicity amplitude arises from scattering n-many particles with a particular configuration of helicities (e.g. two gluons with positive helicity and five with negative helicity).

Spinor–helicity variables make manifest a transformation property of helicity amplitudes: these amplitudes are *little group covariant*. Combined with two further principles—dimensional analysis and locality—this transformation property places strong constraints on the form of possible three-particle scattering amplitudes at first-order in perturbation theory. A system's 'little group' is the subgroup of the Lorentz group that leaves the system's momentum p_μ invariant. For massless particles in four dimensions, p_μ is invariant under two kinds of transformations: translations in space and rotations around the direction of motion. Hence, the little group is the two-dimensional Euclidean group $ISO(2)$, with rotational subgroup $SO(2)$.[24]

Little group transformations leave invariant the 4-momentum of spinor–helicity variables: letting t be a non-zero complex number, $|p\rangle \to t\,|p\rangle$, and $|p] \to t^{-1}\,|p]$. The phases t and t^{-1} cancel, keeping the momentum invariant. However, helicity amplitudes are not invariant under little group scaling. Instead, they are covariant with weight $t_i^{-2h_i}$: $A_n(1^{h_1} \cdots n^{h_n}) = t_i^{-2h_i} A_n(1^{h_1} \cdots n^{h_n})$. Applying little group transformations to helicity amplitudes thereby yields information about the helicity of the scattering particles (Elvang & Huang, 2015, p. 38).

By applying little group covariance alongside a principle of locality, we can use dimensional analysis (counting the mass dimensions of various expressions) to derive the 3-particle helicity amplitudes for massless Yang–Mills theory at tree-level. In the Lagrangian framework, locality imposes the following constraint: the Lagrangian $\mathscr{L}$ must be a function of fields that are (i) defined at a single point and (ii) possess no more than finitely-many derivative terms (with no derivatives in denominators). This locality principle allows us to eliminate some otherwise formally adequate expressions for tree-level 3-particle amplitudes (i.e. expressions that satisfy the constraints from little group covariance and dimensional analysis but violate locality). In this way, one can show that at first-order in perturbation theory, the 3-point amplitudes for massless Yang–Mills theory are fixed entirely by the helicities of the scattering particles (Elvang & Huang, 2015, p. 39). At first-order in perturbation theory, amplitudes with either all positive or all negative helicity particles vanish. If there are two particles of negative helicity and one of positive helicity, then the 3-particle helicity amplitudes are given entirely as a function of angle products:

$$A_3[1^-2^-3^+] = \frac{\langle 12\rangle^3}{\langle 23\rangle\,\langle 31\rangle}; \tag{5}$$

$$A_3[1^-2^+3^-] = \frac{\langle 13\rangle^3}{\langle 12\rangle\,\langle 23\rangle}; \quad A_3[1^+2^-3^-] = \frac{\langle 23\rangle^3}{\langle 12\rangle\,\langle 13\rangle} \tag{6}$$

23 Here, we adopt the notational convenience of letting 'n^{h_n}' denote the kinematic properties of the n-th particle, with momentum 'p_n' and helicity 'h_n'.

24 For a discussion of the little group aimed at philosophers, see Rivat (2023, §3).

Notice that the two negative helicity particles occur in the numerator, while the denominator contains products of the positive helicity particle with both negative helicity particles. If instead there are two positive helicity particles and one negative helicity particle, then the 3-particle amplitudes take exactly the same form, only now expressed in terms of square spinors:

$$A_3[1^+2^+3^-] = \frac{[12]^3}{[23]\,[31]}, \text{ etc.} \tag{7}$$

Using complex analysis, we can develop recursion relations for higher-point tree-level amplitudes as a function of these three-particle amplitudes. In this way, the 3-point amplitudes function as “seed amplitudes” for Yang–Mills theory. Yang–Mills theory is thereby *on-shell constructible*: all tree-level amplitudes can be defined recursively starting from a finite number of seed amplitudes. As the appendix elaborates, deriving recursion relations relies on a third physical principle: *unitarity*. The most familiar definition of unitarity references the scattering matrix. The S-matrix is unitary, meaning that the probabilities of all possible scattering processes always sum to one. Combining unitarity with locality leads to the following factorization property: *at a pole, an amplitude factorizes into left and right subamplitudes, connected by a propagator*. This factorization property underwrites the recursive structure of scattering amplitudes. As discussed in the appendix, these recursion relations require that a pole at infinity vanishes. This typically occurs with felicitous choices about which lines to momentum-shift, but there is no known uniform method for handling the pole at infinity.

Using these recursion relations, the on-shell approach provides epistemic access to the same tree-level amplitudes as Feynman diagrams. At higher-orders in perturbation theory, Feynman diagrams acquire a ‘loop’ structure, with multiple internal propagators connected to each other. On-shell recursion handles these higher-orders through the method of *unitarity cuts* (based on the optical theorem) (Elvang & Huang, 2015, p. 124ff.).[25] In short, one can systematically decompose loop-level amplitudes into products of tree-level amplitudes. Overall then, on-shell recursion shows that knowledge of (i) seed amplitudes, (ii) recursion relations, and (iii) the method of unitarity cuts suffices for knowledge of amplitudes at any order in perturbation theory, for theories that are on-shell constructible. Since none of these calculational techniques, physical principles, or attendant variable changes contradict the method of Feynman diagrams, on-shell recursion provides a compatible problem-solving plan for calculating scattering amplitudes. One can consistently move back and forth between these formulations without contradiction.

6 Some epistemic benefits of on-shell recursion

Having slogged through the rudiments of on-shell recursion, we are entitled to wonder: so what? (at least philosophically speaking). Fortunately, our hard work is amply

25 For more on the generalized unitarity method, see Bern and Huang (2011), Badger et al. (2013b), and Arkani-Hamed et al. (2021, §7).

renumerated. I will argue that the on-shell reformulation of perturbative QFT supports at least three philosophical upshots. First, the derivation of recursion relations provides a particularly compelling example of different epistemic dependence relations between two formulations: by moving to spinor–helicity variables and using complex analysis, we change what we need to know to calculate scattering amplitudes. Second, this reformulation illustrates how it can be worthwhile to identify and utilize basic physical constraints or principles to the full extent possible. While substantial progress in QFT did not require exploiting dimensional analysis, little group scaling, and locality, it is extremely illuminating to discover how much follows from these principles alone. Third, on-shell recursion helps make important patterns manifest, patterns that are otherwise hidden (such as the Parke–Taylor formula). Section 7 develops this third philosophical upshot in detail.

The on-shell formulation teaches us the following: to approximate an n-point scattering amplitude, it suffices to know (i) the lower-point amplitudes and (ii) a relevant set of recursion relations. This knowledge is an instance of an epistemic dependence relation, obtained by reformulating. According to conceptualism, knowledge of these EDRs has epistemic value independently of any practical consequences or methodological benefits. Figuring out what we need to know to solve problems enhances our understanding. Without this knowledge from on-shell recursion, one might erroneously think that they need to know the physical or mathematical content associated with each connected Feynman diagram at a given order in perturbation theory. Even in a context where we choose not to avail ourselves of on-shell recursion, knowledge of this EDR remains valuable. As this case study illustrates, it is epistemically valuable to clarify what we need to know to solve problems, and this is one of the benefits of reformulating. Of course, gaining knowledge of EDRs typically has downstream practical and methodological benefits as well. We can save ourselves a vast amount of computational resources by decomposing higher-point amplitudes into lower-point amplitudes that we have already calculated.

For those who remain skeptical of the non-practical epistemic value of reformulating, some interpretive consequences of on-shell recursion might be more convincing. The EDR behind on-shell recursion supports the following physical interpretation of scattering amplitudes: scattering amplitudes have a recursive structure, wherein part of the physical content of a scattering amplitude is contained within lower-point amplitudes. The success of on-shell recursion provides evidence of this recursive structure in reality (or at any rate, evidence for some physical connection between scattering amplitudes that involve different numbers of particles). Intuitively, knowledge of this recursive structure has epistemic value independently of any practical benefits that might come from exploiting it in calculations. This interpretive point illustrates a philosophical thesis defended by Ruiz de Olano et al. (2022), namely that approximation methods can influence how we interpret a physical theory. Here, we see that a particular formulation of perturbation theory provides evidence for how to interpret the physical content of scattering amplitudes.

On-shell recursion also illustrates a more general moral regarding a particular methodology for reformulating. Schematically, on-shell recursion succeeds by identifying and applying basic physical constraints or principles. The goal is to extract as much physical information from these principles as possible, thereby constraining the form and content of perturbative scattering amplitudes. In particular, we saw how some three-particle amplitudes are completely determined by (i) a principle of locality, (ii) their transformation properties under the little group (stemming from a background principle of Poincaré invariance), and (iii) a logico-physical principle of dimensional analysis.[26] Then, to develop recursion relations, we extract further information from a physical principle of unitarity. Overall, on-shell recursion illustrates how it sometimes pays to 'go back to basics' when reformulating a theory or approximation method.[27]

Nonetheless, it can sometimes take decades to reap the benefits of reformulating a theory using basic principles. In many ways, on-shell recursion is a successful resuscitation of a similar but largely unsuccessful approach from the 1950s–1960s, known as *S*-matrix theory.[28] The *S*-matrix program sought to constrain scattering amplitudes on the basis of a few fundamental physical principles, rather than by postulating Lagrangians for field theories. Perhaps this program floundered because it initially focused on the less tractable case of massive particles. On-shell recursion takes advantage of many of the same physical principles, but it initially focused on massless particles, often in the context of more tractable supersymmetric Yang–Mills theories.[29] Hence, there is no guarantee that this general methodological approach to reformulating will be fruitful.

7 Making properties manifest

Physicists often claim that one formulation or representational choice makes a property manifest, such as a particular symmetry.[30] Perhaps the most common example is 'manifest Lorentz invariance,' wherein a representation makes manifest that an expression is Lorentz invariant. Moreover, by clarifying what we need to know to solve problems, some reformulations make properties or facts *more manifest*, indicating that manifestness admits of degrees. Nonetheless, physicists do not typically specify what it takes for a property to be manifest, as opposed to being hidden or obscured.[31]

Roughly, by calling a property 'manifest,' physicists seem to mean that it is easy or simple to infer that the property obtains. In the case of manifest Lorentz invariance, one

[26] For more on the epistemic benefits of dimensional analysis, see Sterrett (2009) and Lange (2009).

[27] Carcassi and Aidala's *Assumptions of Physics* project (2021) illustrates this methodological moral as well, in the context of classical mechanics, quantum mechanics, and thermodynamics.

[28] For contemporaneous overviews of *S*-matrix theory, see Eden et al. (1966), Olive (1964), and Chew (1966). For a philosophical–historical account, see Cushing (1990).

[29] For discussions of the relationship between the *S*-matrix program and on-shell recursion, see Benincasa and Conde (2012), Feng et al. (2011), and Bern and Huang (2011, p. 2). Illustrating the methodological moral discussed here, Benincasa and Conde note that "ideally one would like to formulate a general *S*-matrix theory starting from a minimal amount of assumptions" (2012, p. 2). This aim goes back to early proponents of the *S*-matrix program (Cushing, 1990, pp. 116–118, 133, 182).

[30] For instance, Cohen, Craig, et al. (2023) discuss how one can make manifest the independence of scattering amplitudes from the choice of field parameterization, an issue I discuss below.

[31] Regarding Yang–Mills theory and perturbative approaches to gravity, physicists sometimes mean something very precise when they say a variable choice 'makes a symmetry manifest,' namely that the variables transform linearly under a given symmetry group. Yet, this definition is too restrictive to capture all uses of 'manifestness.'

can 'read off' Lorentz invariance from index notation, checking that each lower tensor index matches a corresponding upper tensor index. Matching indices indicate that the represented quantity transforms as a Lorentz scalar (and is hence invariant). Yet, what it takes for an inference to be 'easy' or 'simple' is problematically subjective: these factors depend primarily on idiosyncratic features of scientific agents, rather than the formulation or mode of presentation itself. What is easy or simple for one person to infer might be difficult for another, based on differences in their cognitive capacities or skills. Fortunately, from this subjective starting point, we can extract a notion of 'manifestness' that is at least intersubjective. In what follows, I propose and apply a more precise characterization of what it means to make a property manifest.

Lying behind physicists' intuitions about easy or simple inferences are norms governing when an inference is warranted. I propose that a formulation *makes a property manifest* when an agent who understands that formulation is warranted to infer that the property obtains. In other words, the agent's evidence warrants or licenses this inference. Depending on one's preferred philosophical account of epistemic warrant, this makes 'manifestness' at least an intersubjective—if not an entirely objective—feature of a problem-solving formulation. Additionally, by focusing on epistemic warrant, my account avoids appeals to differences in how joint-carving, natural, or fundamental different formulations are. In this way, my proposal is less metaphysically-committed than related proposals regarding *perspicuous representations* (Møller-Nielsen, 2017; North, 2021).

To illustrate my account of making properties manifest, consider again the property of being invariant under Lorentz transformations. As noted above, an expression is manifestly Lorentz invariant provided all lower and upper tensor indices are paired off. Matching indices warrants inferring that the quantity is a Lorentz scalar and hence invariant under Lorentz transformations. For instance, the expression $F_{\mu\nu}F^{\mu\nu}$ is manifestly Lorentz invariant. In contrast, the expression $\int \frac{d^3k}{(2\pi)^3 2w_k}$ is *not* manifestly Lorentz invariant, despite representing a Lorentz invariant measure (here, $w_k := +\sqrt{|k|^2 + m^2}$). Short of performing a calculation, one is not warranted to infer that this expression is Lorentz invariant, simply on the basis of its representational form.

Turning to a slightly more involved example, consider the following Lagrangian density (Cheung, 2017, p. 2):

$$\mathscr{L} = \frac{1}{2}\left[1 + \lambda_1\phi + \frac{1}{2!}\lambda_2\phi^2 + \frac{1}{3!}\lambda_3\phi^3 + \frac{1}{4!}\lambda_4\phi^4 \ldots\right]\partial_\mu\phi\partial^\mu\phi \tag{8}$$

Superficially, this Lagrangian density appears to represent an *interacting* theory, based on the terms involving ϕ^3, ϕ^4, and higher powers of the scalar field ϕ. In fact, this Lagrangian density describes a free, i.e. non-interacting, scalar field. For instance, the non-trivial part of its four-particle scattering amplitude vanishes. Working from the Lagrangian density (8), one could define Feynman rules and compute this four-particle amplitude at first-order in perturbation theory, finding that it equals zero. Likewise, one could calculate the five-particle amplitude, again finding that the contributions from different Feynman diagrams cancel each other out.

As Cheung dramatically notes, "the 14-particle [tree-level] amplitude also vanishes, albeit through the diabolical cancellation of upwards of 5 trillion Feynman diagrams" (2017, p. 3). In saying this, Cheung has not himself written down any of these 5 trillion Feynman diagrams. Instead, it suffices to perform an appropriate *field redefinition* $\phi(x) \rightarrow \tilde{\phi}(x)$ that makes manifest the non-interacting nature of the Lagrangian (8). In general, scattering amplitudes are invariant under local redefinitions (a.k.a. re-parameterizations) of the fields appearing in the Lagrangian density. A redefinition of a field $\phi(x)$ adds any local function $F(\phi(x))$ to that field, provided the derivative of F evaluated at zero equals one (i.e. $F'(0) = 1$). By making an appropriate choice of $F(\phi(x))$, we can transform the Lagrangian density (8) into a manifestly free form, where m is the mass of the field:

$$\tilde{\mathscr{L}} = \frac{1}{2}\partial_\mu \tilde{\phi} \partial^\mu \tilde{\phi} + \frac{1}{2} m^2 \tilde{\phi}^2 \tag{9}$$

Since formulations can make properties *more or less* manifest, my account must allow for gradation. Intuitively, depending on the formulation being used, an agent can be more or less 'close' to having a warranted inference. One way to quantify this notion of closeness relies on the number of epistemically possible solutions at a given stage of problem-solving. A solution counts as epistemically possible provided it has not yet been ruled out by the agent's evidence. Other things equal, an agent is closer to the solution provided fewer non-actual answers are epistemically possible, i.e. provided their evidence has ruled out more non-actual putative solutions. At each stage of calculation, an agent applies an inference rule which changes their evidence (or at least their proximal or salient evidence), restricting which answers are epistemically possible (relative to their proximal evidence).

In a given problem-solving context, each possible solution corresponds to an epistemically possible property that the system might have. The problem is to determine which property or fact obtains, relative to a contrast class of possible properties or facts. With these clarifications in place, I propose the following gradated account of manifestness: compared to a given formulation, another makes a property *more manifest* provided it rules out more epistemically possible properties from the given contrast class. This means that an agent working with the 'more manifest' formulation is closer to being warranted to infer the solution—e.g. the relevant property or fact that actually obtains—compared to another agent who begins with the same initial input information (i.e. background evidence) while using a 'less manifest' formulation.

Comparing the Feynman diagram formulation with on-shell recursion, the latter makes more manifest some striking properties of scattering amplitudes. One paradigmatic example is the Parke–Taylor formula, which characterizes gluon helicity amplitudes in terms of an elegant cyclic structure (provided exactly two of the gluons have negative helicity):

$$A_n[p_1^- p_2^- p_3^+ \cdots p_n^+] = \frac{\langle p_1 p_2 \rangle^4}{\langle p_1 p_2 \rangle \langle p_2 p_3 \rangle \cdots \langle n p_1 \rangle} \tag{10}$$

After calculating the tree-level amplitude for two gluons scattering to four gluons (1986a), Parke and Taylor (1986b) proposed this formula as a well-motivated conjecture. Berends and Giele first proved it in 1988 using off-shell recursion techniques combined with color-ordering and spinor–helicity variables (1988, p. 788ff.). Using Feynman diagrams, a proof of this relation for the special case of a 7-gluon tree amplitude would generically require analyzing 154 diagrams (Elvang & Huang, 2015, p. 34). These calculations would also require choices of field redefinitions and gauge, even though the answer does not depend on these choices. In contrast, on-shell recursion provides a short inductive proof, spanning a mere three pages once the on-shell formulation has been laid out (Elvang & Huang, 2015, pp. 55–57).[32] Intuitively then, an agent using the on-shell formulation is much 'closer' to being warranted to infer the Parke–Taylor formula. In virtue of good variable choices and the physical principles deployed, the on-shell formulation rules out more answers that are initially epistemically possible, regarding the form of these n-point gluon amplitudes. Whereas with the Feynman diagram approach, more epistemically possible but non-actual answers are 'live,' up until the end of the calculation when the Parke–Taylor formula emerges.[33]

As this example suggests, making a property (more) manifest typically provides a number of practical and methodological benefits, such as shortening the length of proofs or calculations. This saves time and resources, making it easier to derive some facts about amplitudes, such as the Parke–Taylor formula. Of deeper philosophical interest, making a property (more) manifest is another way reformulations can provide non-practical epistemic value. Intuitively, it is intellectually valuable to approach logical omniscience. By making a property (more) manifest, a formulation takes us closer to logical omniscience. We spend fewer evidential stages of calculation further away from knowledge of a relevant consequence of our framework or theory.

For another argument that the on-shell reformulation yields something of non-practical epistemic significance, consider the following. The form of the Parke–Taylor formula shows that it does not depend on particular choices of gauge or field redefinition for the quantum fields in the Lagrangian. Intuitively then, we should be able to derive and *understand* the formula without relying on such representational over-specifications.[34] Regardless of any downstream practical benefits, it is intellectually significant to derive this formula without introducing gauge choices and field redefinitions. The on-shell formulation provides a way of achieving this epistemic aim, thereby improving our understanding of the Parke–Taylor formula. In contrast, any way of deriving this formula using Feynman diagrams requires making a particular choice of gauge and field parametrization. Hence, the Feynman diagram formulation is epistemically deficient in a way that the on-shell formulation is not.

[32] See Henn and Plefka (2014, pp. 42–45) for a similar proof.

[33] Henn and Plefka helpfully summarize this difference between the two approaches as follows: "A reason for the complexity of the Feynman diagram calculation is that individual Feynman diagrams are gauge variant and involve off-shell intermediate states in internal propagators. The amplitude, on the other hand, is gauge invariant and only knows about on-shell degrees of freedom. Hence, in going from Feynman diagrams to an amplitude the unphysical degrees of freedom cancel. On-shell approaches that focus on the analytic structure of the final result allow [us] to circumvent these unnecessary complications" (2014, pp. vii–viii).

[34] This desideratum is somewhat analogous to an occasional desire in mathematics for 'purity of proof' or method. For the notion of representational over-specification, see Manders (1999).

Of course, this limitation does not entail that Feynman diagrams obscure *all* physically-relevant properties of scattering amplitudes. As we saw in Section 4, Feynman diagrams make manifest some topological features of scattering amplitudes. For instance, they warrant inferring that any term in the perturbation series which corresponds to a disconnected Feynman diagram does not contribute to the non-trivial part of the scattering amplitude. This is a fact that elementary perturbation theory obscures.

Even compared with the on-shell formulation, there are physical properties that Feynman diagrams succeed at making more manifest. In particular, Feynman diagrams make locality manifest, while on-shell recursion obscures this property within amplitude calculations. To accommodate an inherent gauge redundancy within polarization 4-vectors, the on-shell formalism introduces an arbitrary reference momentum spinor (representing an equivalence class of gauge-related polarization vectors). Problematically, using reference momenta to re-express polarization vectors leads to the appearance of *spurious poles* (Elvang & Huang, 2015, p. 61). Calculating amplitudes using reference momenta frequently leads to expressions of the following form:

$$\frac{1}{[i|\, p_i + \cdots + p_j\, |k\rangle} \tag{11}$$

Interpreted literally, these expressions correspond to non-local interactions, namely a pole in a tree amplitude where a particle fails to go on-shell. They thereby represent an unphysical pole. Ultimately, the residues of these unphysical poles must cancel systematically so that only physically meaningful (i.e. local) poles remain.[35] By contrast, in the Feynman diagram approach—once appropriate gauge choices and field redefinitions are made—locality is manifest throughout the calculation: all poles in a tree amplitude correspond to a propagating particle going on-shell, indicating that all interactions are local.

If one were searching for a grand philosophical moral from these examples it might be this: we should not expect that success at making some properties manifest will translate into success at making *all* physically-relevant or fundamental properties manifest. There may not be a canonically best formalism for making physical properties manifest. This moral is similar to Ruetsche's skepticism that we will ever arrive at a single, unified interpretation of all successful scientific theories (2015, p. 3434). Hence, in at least a small way, this case study provides reasons to doubt that physics will arrive at a uniquely-best language for 'carving nature at its joints.' Fortunately, we have seen how to appraise the non-practical epistemic value of these formulations without relying on putative differences in joint-carving.

8 Conclusions

By examining three formulations of perturbative QFT, I have illustrated numerous ways in which reformulating an approximation method can improve our understanding

[35] Spurious poles arise in the case of handling massive particles as well (Arkani-Hamed et al., 2021, p. 19). For a discussion of how to eliminate spurious poles by changing variables to momentum twistors, see Hodges (2013).

of a physical theory. At each stage of reformulating, we clarify what we need to know to calculate scattering amplitudes. The Feynman diagram formulation shows that we do not need to calculate terms corresponding to disconnected diagrams. The on-shell formulation shows that we typically can construct higher-point amplitudes recursively from a small number of seed amplitudes.

Clearly, these aspects of reformulation provide methodological benefits that are of practical epistemic significance. They facilitate and speed up calculation, while reducing the risk of mistakes for imperfect epistemic agents like ourselves. Beyond these practical benefits, I have defended a stronger claim: these reformulations also provide epistemic benefits of a non-practical sort, thereby supporting my conceptualist account of the 'intellectual significance' of compatible reformulations (2021b, 2023). In particular, I have argued that different formulations of perturbative QFT succeed at making different properties manifest. According to my proposal, a formulation makes a property manifest whenever it warrants inferring that the given property obtains. By gradating this account of manifestness, I captured how some formulations make a property *more* manifest, compared with other formulations. Finally, this case study supports Fraser's (2020) and Ruiz de Olano et al.'s (2022) contention that approximation methods matter for theory interpretation. For instance, on-shell recursion provides evidence that scattering amplitudes possess a recursive structure. Additionally, we have seen how reformulations can reveal patterns that are otherwise hidden or obscured, such as the Parke–Taylor formula for gluon helicity amplitudes.

Appendix

According to a standard principle of locality, all particle interactions are local, represented either by a local field or by an on-shell propagating particle which mediates the interaction (Elvang & Huang, 2015, pp. 44–45). In terms of the analytic structure of tree-level amplitudes, locality posits that *all tree-level poles (with non-vanishing residue) correspond to a propagating particle going on-shell*. Hence, all poles in the amplitude arise from a sum of momenta, P, going on-shell, such that $P^2 = 0$. Where there is a pole, there is a propagator. Unitarity requires that the probability of a scattering process be preserved under insertion of a pole, which requires that the amplitude factorize into a product of left and right subamplitudes: $A_n \xrightarrow[P^2=0]{} A_L \frac{1}{P^2} A_R$. Combined with a simple application of complex analysis, this factorization property results in powerful recursion relations.

Different kinds of recursion relations amount to different ways of formulating the same physical content. These approaches are based on applying Cauchy's residue theorem to an analytic continuation of the amplitude A_n into the complex plane (Elvang & Huang, 2015, Ch. 3). Since we know from locality and unitarity that n-point amplitudes factorize into lower-point amplitudes at poles, we seek to characterize the amplitude in terms of its poles. Cauchy's residue theorem provides one way of expressing a holomorphic function in terms of residues of its poles. This motivates defining a *complex shifted amplitude* $\hat{A}_n(z)$, which is a holomorphic (i.e. complex differentiable) function of a complex variable z. We construct the shifted amplitude $\hat{A}_n(z)$ by shifting some component momenta p_i^μ by an associated complex-valued vector r_i^μ times z. The shift vectors r_i^μ are constrained such that the shifted momenta $\hat{p}_i^\mu(z)$ remain on-shell and satisfy momentum conservation.[36] When $z = 0$, the shifted momenta return to being wholly real-valued, and we recover the unshifted amplitude of interest $A_n = \hat{A}_n(0)$. $\hat{A}_n(z)$ is thus the analytic continuation of A_n into the complex plane. More usefully, A_n is the residue at $z = 0$ of the holomorphic function $\frac{\hat{A}_n(z)}{z}$.[37]

Since $\frac{\hat{A}_n(z)}{z}$ is holomorphic except at finitely-many isolated singularities (i.e. its poles), we can apply Cauchy's residue theorem to calculate A_n as a sum of residues. First, note that $\hat{A}_n(0) = \frac{1}{2\pi i} \oint\limits_{C_0} \frac{\hat{A}_n(z)}{z} \mathrm{d}z$, where C_0 is a closed contour around the pole at $z = 0$ (and containing no other singularities). We proceed to deform the contour C_0 to oppositely oriented contours surrounding the other poles, including possibly a pole at infinity. Then, applying Cauchy's residue theorem, we see that $\hat{A}_n(0) = B_\infty - \sum_{z_I} Res_{z=z_I}\left[\frac{\hat{A}_n(z)}{z}\right]$. Here, '$B_\infty$' denotes the residue of the pole at infinity. As mentioned in Section 5, the construction of on-shell recursion relations requires showing that B_∞ equals zero. To summarize what we've done so far: by transforming the on-shell amplitude into a holomorphic function of a complex variable, we can determine its value by focusing on the residues of poles of this holomorphic function. Cauchy's theorem lets us express the physical, unshifted amplitude as a summation of residues of the shifted amplitude divided by z.

Our goal then is to calculate these residues $Res_{z=z_I}\left[\frac{\hat{A}_n(z)}{z}\right]$ of the various poles. As it turns out, each residue equals a product of lower-point amplitudes connected by a propagator. This is how we finally arrive at recursion relations for higher-point amplitudes in terms of lower-point amplitudes. First, we define the following sums of momenta and shift vectors: $P_I = \sum_{i \in I} p_i$, $\hat{P}_I = \sum_{i \in I} \hat{p}_i$, and $R_I = \sum_{i \in I} r_i$. Using the constraints on the shift vectors r_i^μ, we can show that $\hat{P}_I^2 = P_I^2 + z \cdot 2P_I \cdot R_I$. We want to factor this expression so that when $\hat{P}_I^2 = 0$ (i.e. when $1/\hat{P}_I^2$ has a pole), we have a zero at $z = z_I$. This lets us define the other poles z_I corresponding to the sums of shifted momenta equaling zero. A short manipulation shows that we should define z_I as $\frac{-P_I^2}{2P_I \cdot R_I}$. Taking the reciprocal of $\hat{P}_I^2$, we then have $\frac{1}{\hat{P}_I^2} = \frac{1}{z - z_I} \frac{-z_I}{P_I^2}$, as desired.

By this construction, each simple pole z_I corresponds to a sum of complex shifted momenta $\hat{P}_I = \hat{p}_1 + \cdots + \hat{p}_k$ going to zero. By locality, we know that each of these poles corresponds to an on-shell propagator term $1/\hat{P}_I^2$. Then, by unitarity, we know that we can factorize the shifted amplitude as $z \to z_I$:

$$lim_{z \to z_I} \hat{A}_n(z) = \hat{A}_L(z_I) \frac{1}{\hat{P}_I^2} \hat{A}_R(z_I) \tag{12}$$

[36] Specifically, we require that $\sum_{i=1}^n r_i^\mu = 0$; $r_i \cdot r_j = 0 \ \forall i, j$; and $p_i \cdot r_i = 0 \ \forall i$ (Elvang & Huang, 2015, p. 50).

[37] To see this, recall from elementary complex analysis that the residue of a simple pole z_I is given by the limit as $z \to z_I$ of the holomorphic function times $(z - z_I)$. In this case, $z = 0$ is a simple pole of $\frac{\hat{A}_n(z)}{z}$. Hence, we have $Res(\frac{\hat{A}_n(z)}{z},\ z = 0) = lim_{z \to 0}(z - 0)\ \frac{\hat{A}_n(z)}{z} = lim_{z \to 0}\ \hat{A}_n(z) = A_n$.

This expresses the amplitude in terms of the complex shifted momenta $\hat{P}_I^2$, but ultimately we want a factorization using the real, physical momenta P_I instead. Our goal remains to calculate the residues $Res_{z=z_I}\left[\frac{\hat{A}_n(z)}{z}\right]$. By a simple algebraic manipulation, we can show that $Res_{z=z_I}\left[\frac{\hat{A}_n(z)}{z}\right]$ actually equals $-\hat{A}_L(z_I)\frac{1}{P_I^2}\hat{A}_R(z_I)$.[38] Hence, substituting this expression into our earlier expression from Cauchy's residue theorem, we discover the following recursion relation:

$$A_n = \hat{A}_n(0) = B_\infty + \sum_{diagrams_I}\left[\hat{A}_L(z_I)\frac{1}{P_I^2}\hat{A}_R(z_I)\right] \tag{13}$$

This represents our amplitude of interest as a product of two shifted amplitudes $\hat{A}_L(z_I)$ and $\hat{A}_R(z_I)$ (corresponding to left and right parts of an on-shell diagram), with a propagator term $1/\hat{P}_I^2$ in between. Note that the left and right subamplitudes ($\hat{A}_L$ and $\hat{A}_R$) must each involve fewer than n-particles since they each have at least one external particle leg. Hence, provided that $B_\infty = 0$, we have just expressed our n-point amplitude of interest, A_n, as a summation over products of lower-point (complex-shifted) on-shell amplitudes $\hat{A}_L$ and $\hat{A}_R$. We could apply the same recursion relation to $\hat{A}_L$ and $\hat{A}_R$ to express them as functions of lower-point amplitudes, all the way until we reach the seed amplitudes (assuming the theory is on-shell constructible).

For this method to succeed, the residue of the pole at infinity (B_∞) must equal zero. This means that it is necessary to know the behavior of the amplitude at large z, i.e. at the boundary of the complex plane. There is no known systematic way of calculating this for arbitrary quantum field theories (Elvang & Huang, 2015, p. 52). It generally must be shown on a case-by-case basis. Furthermore, only certain combinations of shifted external momenta succeed at satisfying this boundary condition. For instance, some choices of shifts in pure Yang–Mills theory do not lead to a vanishing residue at infinity. This boundary condition problem provides a good example of a *lack* of uniformity of treatment, since we don't have a uniform method for calculating this residue systematically. In contrast, approaches based on Feynman diagrams do not encounter a similar problem, since they do not have to calculate this residue for a pole at infinity.

One of the most popular on-shell methods uses *BCFW recursion relations*, named after Ruth Britto, Freddy Cachazo, Bo Feng, and Edward Witten (2005). The resulting on-shell diagrams are known as BCFW diagrams or BCFW terms/subamplitudes. The BCFW recursion relations make the generic recursion relations described above more convenient by shifting the fewest possible external momenta. This decreases the number of on-shell diagrams that need to be considered to compute an n-point amplitude.

[38] This follows from the fact that $1/\hat{P}_I^2$ equals $\frac{-z_I}{z-z_I}\frac{1}{P_I^2}$. Then, since z_I is a simple pole, we have $Res_{z=z_I}\left[\frac{\hat{A}_n(z)}{z}\right] = lim_{z\to z_I}(z - z_I)\frac{\hat{A}_n(z)}{z} = lim_{z\to z_I}\frac{z-z_I}{z}\hat{A}_n(z) = lim_{z\to z_I}\frac{z-z_I}{z}\hat{A}_L(z_I)\frac{1}{\hat{P}_I^2}\hat{A}_R(z_I)$. Then, substituting for $\hat{P}_I$, we see that this equals $lim_{z\to z_I}\left[\frac{z-z_I}{z}\frac{-z_I}{z-z_I}\hat{A}_L(z_I)\frac{1}{P_I^2}\hat{A}_R(z_I)\right] = -\hat{A}_L(z_I)\frac{1}{P_I^2}\hat{A}_R(z_I)$.

Regarding their scope, BCFW recursion has been shown to work in a variety of different quantum field theories and gravity theories. These include gluon scattering in Yang–Mills (YM) theory, scalar-QED, $\mathcal{N} = 4$ super Yang–Mills theory[39], and graviton scattering at tree-level (Elvang & Huang, 2015, pp. 62-65). In pure YM theory, $\mathcal{N} = 4$ super YM theory, and gravity scattering, all tree-level amplitude information is ultimately contained in 3-point amplitudes. This suggests interpreting the higher-point interaction vertices in the Lagrangians for these theories as not carrying additional physical content beyond the on-shell three-point interactions. For instance, the four-point interaction term in the Feynman diagram formulation of pure YM theory may merely serve to maintain the gauge invariance of the Lagrangian (2015, p. 62). It is an open question whether BCFW recursion extends to scattering arbitrary massive particles (Arkani-Hamed et al., 2021, pp. 65–66), although it has been shown to apply in some specific cases (Badger et al., 2006; Ozeren & Stirling, 2006).

For arbitrary renormalizable quantum field theories in four dimensions, BCFW recursion does not generally work. However, other recursion methods have been developed that succeed. In particular, by shifting *all* of the external momenta lines, one can develop recursion relations using three-point and four-point scattering amplitudes as the seed amplitudes (Cheung, 2017, pp. 27–29; Cohen et al., 2011). This *all-line recursion* is computationally more complicated than BCFW recursion. Since more momenta are shifted, there are more subamplitudes to consider when factorizing. Nonetheless, it has been shown that if the theory contains only massless particles, it is necessary to shift only five external lines, and in many cases three lines suffice (Cheung et al., 2015). Additionally, on-shell recursion succeeds even for some non-renormalizable theories with massless particles, including those that do not have derivative interactions (Cheung et al., 2015). More recently, on-shell recursion methods have been extended to the Standard Model, viewed as an effective field theory (Aoude & Machado, 2019; Huber & De Angelis, 2021; Liu et al., 2023). Constructing recursion relations for theories with massive particles is an active area of research (Ballav & Manna, 2022).

Acknowledgements For feedback on versions of this material, I thank Dave Baker, Gordon Belot, Henriette Elvang, and Laura Ruetsche. For questions, comments, and suggestions, I thank audiences at the University of Pittsburgh Early Career Workshop in Philosophy and History of Physics, the University of Michigan workshop on reformulations, the Approximations to Second Order workshop at the Max Planck Institute of the History of Science, and Foundations 2023 at the University of Bristol, especially Alexander Blum, Arianna Borrelli, Jeremy Butterfield, James Fraser, Henrique Gomes, Mike Miller, Sébastian Rivat, James Wells, and Porter Williams. Thanks to Pablo Ruiz de Olano for organizing the workshop that crystallized this project and to Ken Manders for inspiring it.

[39] $\mathcal{N}$ counts the number of supercharges Q^A and $Q^\dagger_A$, which are supersymmetry generators that transform bosons into fermions and vice versa.

References

Abreu, S., Dormans, J., Cordero, F. F., Ita, H., Page, B., & Sotnikov, V. (2019). Analytic form of the planar two-loop five-Parton scattering amplitudes in QCD. *Journal of High Energy Physics, 2019*(84), 1–31.

Aoude, R., & Machado, C. S. (2019). The rise of SMEFT on-shell amplitudes. *Journal of High Energy Physics, 2019*(58), 1–24.

Aoyama, T., Hayakawa, M., Kinoshita, T., & Nio, M. (2012). Tenth-order QED Lepton anomalous magnetic moment: Eighth-order vertices containing a second-order vacuum polarization. *Physical Review D, 85.033007*(3), 1–17.

Arkani-Hamed, N., Huang, T.-C., & Huang, Y.-T. (2021). Scattering amplitudes for all masses and spins. *Journal of High Energy Physics, 2021*(70), 1–75.

Badger, S. (2016). Automating QCD amplitudes with on-shell methods. *Journal of Physics: Conference Series, 762*(1), 012057.

Badger, S., Biedermann, B., Hackl, L., Plefka, J., Schuster, T., & Uwer, P. (2013). Comparing efficient computation methods for massless QCD tree amplitudes: Closed analytic formulas versus Berends-Giele recursion. *Physical Review D*. 10.1103/PhysRevD.87.034011

Badger, S., Frellesvig, H., & Zhang, Y. (2013b). A two-loop five-gluon helicity amplitude in QCD. *Journal of High Energy Physics, 2013*(45), 1–31.

Badger, S., Glover, E., & Khoze, V. V. (2006). Recursion relations for gauge theory amplitudes with massive vector bosons and fermions. *Journal of High Energy Physics*.

Ballav, S., & Manna, A. (2022). Recursion relations for scattering amplitudes with massive particles II: Massive vector bosons. *Nuclear Physics B, 983*, 115935.

Batterman, R. W. (2002). *The devil in the details: Asymptotic reasoning in explanation, reduction, and emergence*. Oxford University Press.

Benincasa, P., & Conde, E. (2012). Exploring the S matrix of massless particles. *Physical Review D, 86*, 025007.

Berends, F. A., & Giele, W. T. (1987). The six-gluon process as an example of Weyl–van der Waerden spinor calculus. *Nuclear Physics B, 294*, 700–732.

Berends, F. A., & Giele, W. T. (1988). Recursive calculations for processes with n gluons. *Nuclear Physics B, 306*(4), 759–808.

Bern, Z., & Huang, Y.-T. (2011). Basics of generalized unitarity. *Journal of Physics A: Mathematical and Theoretical, 44*(45), 1–32.

Britto, R., Cachazo, F., Feng, B., & Witten, E. (2005). Direct proof of the tree-level scattering amplitude recursion relation in Yang–Mills theory. *Physical Review Letters, 94*(18), 181602.

Campbell, J. M., & Ellis, K. R. (2023). Top tree amplitudes for higher order calculations. *Journal of High Energy Physics, 2023*(125), 1–41.

Carcassi, G., & Aidala, C. A. (2021). *Assumptions of physics*. Michigan Publishing.

Cheung, C. (2017). TASI lectures on scattering amplitudes. arXiv Preprint. arXiv:1708.03872

Cheung, C., Shen, C.-H., & Trnka, J. (2015). Simple recursion relations for general field theories. *Journal of High Energy Physics, 2015*(6), 118.

Chew, G. F. (1966). *The analytic S-matrix: A basis for nuclear democracy*. W. A. Benjamin.

Cohen, T., Craig, N., Lu, X., & Sutherland, D. (2023). On-shell covariance of quantum field theory amplitudes. *Physical Review Letters, 130*, 041603.

Cohen, T., Elvang, H., & Kiermaier, M. (2011). On-shell constructibility of tree amplitudes in general field theories. *Journal of High Energy Physics, 2011*(4), 53.

Cordero, F. F., von Manteuffel, A., & Neumann, T. (2022). Computational challenges for multi-loop collider phenomenology: A Snowmass 2021 white paper. *Computing and Software for Big Science, 6*(14), 1–14.

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

Dixon, L. J. (2016). A brief introduction to modern amplitude methods. In *Journeys through the precision frontier: Amplitudes for colliders: TASI 2014 Proceedings of the 2014 theoretical advanced study institute in elementary particle physics* (pp. 39–97). World Scientific.

Eden, R. J., Landshoff, P. V., Olive, D. I., & Polkinghorne, J. C. (1966). *The analytic S-matrix*. Cambridge University Press.

Elvang, H., & Huang, Y.-T. (2015). *Scattering amplitudes in gauge theory and gravity*. Cambridge University Press.

Feng, B., Huang, R., & Jia, Y. (2011). Gauge amplitude identities by on-shell recursion relation in S-matrix program. *Physics Letters B, 695*(1), 350–353.
Fillion, N. (2021). Semantic layering and the success of mathematical sciences. *European Journal for Philosophy of Science, 11*(91), 1–25.
Fletcher, S. C. (2019). Minimal approximations and Norton's dome. *Synthese, 196*(5), 1749–1760.
Fraser, J. D. (2020). The real problem with perturbative quantum field theory. *British Journal for the Philosophy of Science, 71*(2), 391–413.
Gastmans, R., & Wu, T. T. (1990). *The ubiquitous photon: Helicity methods for QED and QCD*. Clarendon Press.
Gunson, J. (1965). Unitarity and on-mass-shell analyticity as a basis for S-matrix theories: I. *Journal of Mathematical Physics, 6*(6), 827–844.
Heinrich, G. (2021). Collider physics at the precision frontier. *Physics Reports, 922*, 1–69.
Henn, J. M., & Plefka, J. C. (2014). *Scattering amplitudes in gauge theories* (Vol. 883). Springer.
Hodges, A. (2013). Eliminating spurious poles from gauge-theoretic amplitudes. *Journal of High Energy Physics, 2013*(135), 1–22.
Houriet, A., & Kind, A. (1949). Classification invariante des termes de la matrice S. *Helvetica Physica Acta, 22*(3), 319–330.
Huber, M. A., & De Angelis, S. (2021). Standard model EFTs via on-shell methods. *Journal of High Energy Physics, 2021*(221), 1–54.
Hunt, J. (2021a). Interpreting the Wigner–Eckart theorem. *Studies in History and Philosophy of Science, 87*, 28–43.
Hunt, J. (2021b). Understanding and equivalent reformulations. *Philosophy of Science, 88*(5), 810–823.
Hunt, J. (2023). Epistemic dependence and understanding: Reformulating through symmetry. *British Journal for the Philosophy of Science, 74*(4), 941–974.
Huss, A., Huston, J., Jones, S., & Pellen, M. (2023). Les Houches 2021-Physics at TeV colliders: Report on the standard model precision wishlist. *Journal of Physics G: Nuclear and Particle Physics, 50*, 043001.
Kaiser, D. (2005). *Drawing theories apart: The dispersion of Feynman diagrams in postwar physics*. University of Chicago Press.
Kinoshita, T. (1989). Electron g-2 and high precision determination of α. In G. F. Bassani, M. Inguscio, & T. W. Hänsch (Eds.), *The hydrogen atom* (pp. 247–256). Springer.
Lancaster, T., & Blundell, S. J. (2014). *Quantum field theory for the gifted amateur*. Oxford University Press.
Lange, M. (2009). Dimensional explanations. *Noûs, 43*(4), 742–775.
Liu, H., Ma, T., Shadmi, Y., & Waterbury, M. (2023). An EFT Hunter's guide to two-to-two scattering: HEFT and SMEFT on-shell amplitudes. *Journal of High Energy Physics, 2023*(241), 1–31.
Manders, K. (1999). Euler or Descartes? Representation and responsiveness. Unpublished Manuscript.
Miller, M. E. (2023). Mathematical structure and empirical content. *British Journal for the Philosophy of Science, 74*(2), 511–532.
Møller-Nielsen, T. (2017). Invariance, interpretation, and motivation. *Philosophy of Science, 84*, 1253–1264.
Morrison, M. (2015). *Reconstructing reality: Models, mathematics, and simulations*. Oxford University Press.
North, J. (2021). *Physics, structure, and reality*. Oxford University Press.
Ochirov, A. (2018). Helicity amplitudes for QCD with massive quarks. *Journal of High Energy Physics, 2018*(89), 1–21.
Olive, D. I. (1964). Exploration of S-matrix theory. *Physical Review, 135*(3B), B745–B760.
Ozeren, K., & Stirling, W. (2006). Scattering amplitudes with massive fermions using BCFW recursion. *The European Physical Journal C, 48*, 159–168.
Parke, S. J., & Taylor, T. R. (1986a). The cross section for four-gluon production by gluon-gluon fusion. *Nuclear Physics B, 269*, 410–420.
Parke, S. J., & Taylor, T. R. (1986b). Amplitude for n-gluon scattering. *Physical Review Letters, 56*(23), 2459–2460.
Rivat, S. (2021). Effective theories and infinite idealizations: A challenge for scientific realism. *Synthese, 198*, 12107–12136.
Rivat, S. (2023). Wait, why gauge? *British Journal for the Philosophy of Science*. https://doi.org/10.1086/727736
Ruetsche, L. (2015). The Shaky Game +25, or: On locavoracity. *Synthese, 192*, 3425–3442.

Ruiz de Olano, P., Fraser, J. D., Gaudenzi, R., & Blum, A. S. (2022). Taking approximations seriously: The cases of the Chew and Nambu-Jona-Lasinio models. *Studies in History and Philosophy of Science, 93*, 82–95.
Schuster, T. (2014). Color ordering in QCD. *Physical Review D, 89*, 105022.
Srednicki, M. (2007). *Quantum field theory*. Cambridge University Press.
Sterrett, S. G. (2009). Similarity and dimensional analysis. In A. Meijers (Ed.), *Philosophy of technology and engineering sciences* (pp. 799–823). Elsevier/North-Holland.
Wick, G. C. (1950). The evaluation of the collision matrix. *Physical Review, 80*(2), 268–272.
Wüthrich, A. (2010). *The genesis of Feynman diagrams*. Springer.