Effective and selective realisms*

John Dougherty†

Abstract

Scientific realists argue that empirically successful theories latch on to unobservable features of reality. But it is often thought that conventional theories of particle physics do not deserve realist commitment, despite their outstanding empirical success. Recently, a number of "effective" realisms have argued that we should distinguish between the low- and high-energy claims of particle theory and that we can and should be realist about the former but not the latter. I present a *reductio ad absurdum* against the most naive extension of this proposal to the most empirically successful theories of particle physics, such as quantum electrodynamics. By considering two replies to this argument, I distinguish two forms of effective realism. A conservative form hews closely to traditional forms of realism, and the resources of this tradition allow conservative effective realism to avoid the *reductio*; however, this form of effective realism is left without a positive account of quantum electrodynamics. A more radical form of effective realism can account for quantum electrodynamics, but it requires substantial development, along with a revision of the terms of the realist debate.

1 Introduction

Scientific realists are in a tough spot. They advise positive epistemic attitudes toward our best scientific theories, and they justify this advice by appeal to these theories' achievements—the truth of these theories is the best explanation for their success, the realist might say (Boyd, 1989; Psillos, 2005; Putnam, 1975). This would all be well and good were it not for quantum field theories (QFTs). These theories boast some of history's most dazzling successes in the face of precision tests, so we should be realist about them if we should be realist about any theory. But they rebuke most familiar positive epistemic attitudes, like belief in the truth of what the theories say or in the successful reference of their terms for unobservables; indeed, it's been argued on these grounds that QFTs are not

*This is the author's accepted manuscript version of a paper forthcoming in *The British Journal for the Philosophy of Science* (`https://doi.org/10.1086/724978`).

†I am grateful for helpful discussion and feedback from Kathleen Connelly, Erik Curiel, Johannes Kleiner, Leonardo Moauro, Daniele Oriti, Jer Steeger, Shawn Tinghao Wang, Porter Williams, two anonymous reviewers, and audiences at the MCMP, the Lichtenberg group in History and Philosophy of Physics, and EPSA2021.

even eligible for realist commitment and that realists must wait for some future alternative to QFTs (Fraser, 2009, 2011; Halvorson and Müger, 2007; Kuhlmann, 2010). So the recent philosophical literature comes as a relief when it offers a new kind of "effective" realism that's well adapted to the empirically successful quantum field theories found in the laboratory (Fraser, 2018, 2020a,b; Hancox-Li, 2015; Rivat, 2019, 2021; Rueger, 1990; Ruetsche, 2018, 2020; Williams, 2015, 2019, 2021).

In this paper I distinguish two versions of effective realism and argue that one is more promising—though less developed—than the other. As I explain in Section 2, the basic insight of effective realism is that the philosophically worrisome features of a QFT can be relegated to extremely high energies, well beyond the theory's domain of applicability. Once the high-energy vices are separated out from the low-energy virtues, the realist can ignore the former and commit themselves to the latter—that is, explain the theory's (low-energy) empirical successes by postulating its (low-energy) descriptive accuracy and thereby giving some reason to expect its continued (low-energy) success in the future. My aim to show that this basic insight can be fleshed out in at least two ways which ought to be distinguished, for it makes a practical difference.

Section 3 sets up a *reductio* as a framing conceit for pursuing this aim. The most naive and direct implementation of the effective realist proposal does not work for many QFTs; in particular, it does not extend to quantum electrodynamics (QED), the quantum field theory with the most dramatic empirical success, or to other gauge theories. In these theories, the story that effective realists usually tell fails to separate the low- and high-energy descriptions, and philosophical problems remain at low energies. This *reductio* is a framing conceit based on a naive and direct rendering of the effective realist story in these extended domains. That is, no one could reasonably be said hold the oversimplified premises that generate the contradiction, and so it is not intended as a serious argument against effective realism. It is a straw man set up so that we may consider different ways of knocking it down. It serves to distinguish two versions of effective realism by distinguishing their reasons for rejecting the *reductio*.

The rest of the paper considers these two natural replies on the effective realist's behalf. These replies are based on two incompatible reasons to think that Section 3 incorrectly extends the effective realist proposal to QED. These different reasons are motivated by two different attitudes toward interpretation found in the effective realist movement. Section 4 describes a conservative tendency of effective realism, which assimilates it to other currently popular forms of realism. In particular, the conservative effective realist endorses the broadly logicist thesis that mathematical objects are enough like models of a logical system that they admit literal interpretations and the semantic thesis that they must be so interpreted. Given these semantic commitments, the conservative effective realist can block the argument of Section 3. However, as I explain in Section 5, they are left without a positive alternative. At best they can appeal to a yet-to-be-developed alternative mathematical reconstruction of QFTs; that is, they are in the same position as their ineffective predecessors.

Sections 6 and 7 consider a more radical reading of effective realism on which effective realists should reject the notion of "what a theory literally says" that's adopted by traditional realisms. The semantic tradition of the conservative effective realist is inflexible: once the theory's mathematical formalism is fixed, so is its semantic content; our only remaining choice is which of these contents to believe. The radical reading of effective realism that I offer below drops the conservative's logicist policy on mathematics, replacing it with case-by-case analyses of the semantic content of mathematical representations. This radical is still a realist, and they still endorse a literal interpretation of the semantic content of a given theory, but they disagree with the conservative on how we are to extract semantic content from a theory's mathematics. The radical allows interpretations of a theory's mathematics that explicitly contradict the interpretation forced on the conservative, and may thereby go beyond the conservative's mere withholding of endorsement from this or that claim. This affords them a second method of rejecting the *reductio* of Section 3. The semantic underpinnings of radical effective realism are underdeveloped; for instance, it is not obvious that the outcomes of their interpretive policies are sufficiently unique to satisfy the realist's semantic demands. Nevertheless, I think this case shows that the radical's position deserves further attention.

2 Realism and cutoffs

Effective realism is specifically designed to permit realism about QFTs, our best theories of particle physics.[1] QFTs have historically posed a problem for realism: realists aim to explain a theory's empirical success by appealing to its truth, but QFTs seem insufficiently rigorous, inconsistent, and *ad hoc* to be true (Fraser, 2020a, 392). Effective realists argue that in any QFT these problems can be limited to high energies, beyond the theory's domain of applicability, and they offer a realism about the theory's less problematic description of the low energy world to which it applies.

Like some other versions of scientific realism, effective realism is meant to be an explanation (Fraser, 2020b, 276). As Putnam memorably put it, the thought is that realism "is the only philosophy that doesn't make the success of science a miracle" (1975, 73). Our best scientific theories are remarkably competent at describing, predicting, and explaining a wide range of empirical phenomena. According to the realist, these theories succeed at these tasks by correctly describing the unobservable causes of the target phenomena. This proposal has a metaphysical payoff: because our theories latch on to the causes responsible for observed regularities, they provide reason to expect these regularities in future phenomena. Explanationist realism promises robust philosophical results given the observed success of science and a seemingly thin observation about what needs explanation.

The realist demand for explanation should be felt most keenly when it comes to QFTs, our current best theories of particle physics. In particular, QED—the

[1] Except in Footnote 10, I focus exclusively on perturbative QFTs.

standard quantum-field-theoretic model of the electromagnetic interaction—has been verified by some of history's most precise tests, with predictions correct to one part in ten billion (Hanneke et al., 2008). On the realist's explanationist intuition, a scientific theory succeeds by providing a correct description of unobservable reality. The dramatic successes of QFTs like QED and the Standard Model of particle physics must then mean that their descriptions of the unobservable are substantially correct. We should be realist about them if we're realist about anything.

However, the specter of inconsistency has scared philosophers away from taking QFTs as serious candidates for realist commitment. A correct description of the unobservable must at least be consistent. But it's not obvious that the mathematical manipulations appearing in quantum-field-theoretic arguments can be placed in a well defined mathematical setting: "it is not so clear where 'QFT' can be located in the mathematical universe" (Halvorson and Müger, 2007, 731). As such, we have no assurance that QFT admits a consistent interpretation at all. So, despite its dazzling successes, the conventional framework of QFT is often taken to be a remarkably effective tool for predicting the outcomes of particle experiments with little further content (Fraser, 2009, 2011; Kuhlmann, 2010).

The most notorious threats of inconsistency in QFT are the divergences encountered in nearly every calculation. For instance, consider an experiment in which we attempt to collide two electrons. A QFT provides an inventory of basic interaction types, each with a particular weight, where the weight is a function of quantities like the masses and charges of the interacting particles. In QED, for example, an electron or positron can gain or lose energy by absorbing or emitting a photon, a photon can decay into an electron and a positron, splitting its energy between the two resulting particles, and an electron and a positron can annihilate to produce a photon with their combined energy. The probability that exactly two electrons will exit our scattering experiment is given by the weighted sum over all sequences of basic interactions that begin and end with two electrons. The probability of a final state with two electrons and a photon is the sum over all sequences ending with two electrons and a photon, and so on. At least, this is the heuristic. The problem, for the physicist and realist alike, is that all of these infinite sums diverge. The probability density for every experimental outcome is infinite, so outcomes cannot be consistently assigned probabilities by taking ratios of densities. In particular, the theory is not unitary: the sum of probabilities over all outcomes is not one. So the theory is inconsistent and makes no predictions, if read at face value.

Physicists avoid these divergences by modifying the theory, which makes the case for realism even harder. To avoid divergences from infinite sums, you can just sum over finitely many terms, ignoring the rest. In the simplest situations, you can do this by choosing a large but arbitrary energy Λ and summing over only those basic interactions that take place at energies below Λ. Since all the sums in the modified theory are finite, we may consistently assign normalized probabilities to outcomes of scattering experiments to give a unitary theory. These probabilities will be functions of the energy scale Λ and physical quantities like masses and charges. To obtain more precise predictions, choose a larger cutoff Λ' and renormalize the weights on all the basic interaction types so that probabilities for scattering outcomes sum to one. It's these probabilities that QED predicts with exceeding accuracy. But this means that the empirical successes of QED are obtained by explicitly ignoring the theory's description of some unobservable particle interactions. So realism about these descriptions can't be used to explain the theory's success.

Inconsistent treatment of high-energy details is an obstacle for realism, but not a dead end. Realists should not want to commit to every last detail of a successful theory's description of the unobservable. One reason is that science isn't over: the future will bring new theories for the realist to be realist about, and historical considerations suggest that these future theories will radically differ from today's, so we should think that today's theories are substantially false in some details (Hesse, 1976; Laudan, 1981; Stanford, 2006). A selective realist tries to solve this problem by restricting their commitment to particular features of a theory—structural features, say, or features that underwrite novel predictions—and taking the correct description of these features to explain the theory's success (Worrall, 1989; Psillos, 2005; Chakravartty, 2007). Since the details of the description above the cutoff Λ are ignored when making predictions, the effective realist proposes that we reserve our ontological commitment for the theory's description of phenomena at scales well below Λ.

Effective realism's central insight is that the physicist's justifications for ignoring details of the theory's high-energy description dovetail with the philosopher's reasons for being selective about their loci of realist commitment. This joinery is accomplished by studying how the theory's description of the phenomena depends on the choice of energy scale. The framework for this study is the effective field theory perspective, on which a QFT is taken to be well defined only below some high energy scale Λ at which it breaks down and gives inconsistent results. Because Λ is large compared to the theory's scale of applicability ϵ, the ratio ϵ/Λ can be treated as zero for all practical purposes. Working under this assumption, one can study how the theory's predictions depend on ϵ/Λ. If some quantity depends only on positive powers of ϵ/Λ, then these can consistently be set to zero, making the predicted value effectively independent of Λ. Predictions that ignore details above a cutoff Λ are therefore justified if they only depend on positive powers of ϵ/Λ, because in this case it makes no practical difference to raise the cutoff to a higher value.

This effective field theory perspective has many advantages for the realist (Fraser, 2018, 2020b; Williams, 2019). For our purposes, the most important is that effective field theory methods afford a criterion of positive realist commitment with a natural connection to the selective strategy for avoiding the pessimistic induction. This criterion is based on a conjecture about the relationship between an effective field theory's descriptions at the scales ϵ and Λ. Following the work of Wilson and others in the 1960s and 1970s, the theory's description at scale ϵ can be interpreted as a coarse-graining of the description at Λ (Wilson and Kogut, 1974). At the scale Λ, the basic interaction of an electron emitting a photon can be understood as a black box encoding the detailed high-energy interactions that begin with a single electron and end with an electron and a photon. As the energy scale is lowered, more and more processes are added to this black box, and the weight for this basic interaction is correspondingly readjusted. It turns out that for some families of theories the weights on all but finitely many basic interactions will be multiplied by ϵ/Λ in this procedure, and may therefore be neglected (Polchinski, 1984). In other words, no matter what these theories say about physics at scale Λ, they will all qualitatively agree at low energies, disagreeing only about the values of finitely many physical quantities. So the realist can rest assured that the qualitative low-energy features of today's successful QFTs will not be revised in future theories, since these future theories can differ only in the values of certain constants.

3 The problem

The effective realist's coarse-graining picture presupposes that a QFT's low- and high-energy descriptions may be disentangled and that inconsistencies in the latter can be effectively quarantined behind a cutoff Λ. On the most naive and flat-footed extension of the usual effective realist story to other examples, this presupposition often fails. This section gives a precise illustration of the failure in terms of nonunitarity of QED. Following the effective realist's suggestion in this naive way leads to violations of basic assumptions of the theory—violations that appear at all energy scales, not just very high ones. So it does not solve the realist's problem with QFT. Of course, the effective realist will rightly object that the argument in this section incorrectly extends effective realism to these theories. But different ways of rejecting this *reductio* point to different tendencies in the effective realist movement and, correspondingly, different alternative positive stories.

Effective realism is designed to account for the empirical successes of QFTs—paradigmatically, predictions concerning scattering probabilities. Scattering interactions send an initial state $|i\rangle$ to a final state $|f\rangle$ with amplitude $\langle f|\mathcal{M}|i\rangle$, up to a conventional normalization. Orthodox quantum theory assumes that the scatterings encoded by $\mathcal{M}$ are unitary, so that they preserve the Hilbert space norm and, therefore, total probability. If a proposed scattering operator were nonunitary, it could predict the sum of probabilities of outcomes for some experiment to be other than 100%, contradicting the assumptions of probability theory.

Effective realism is, in part, an interpretation of the mathematical manipulations used to calculate the operator $\mathcal{M}$. Heuristically, the amplitude can be expressed in terms of integrals of classical data. To take an example that's common in the effective realism literature, consider the theory of a quartically self-interacting massive real scalar field ϕ. The scattering operator in this theory is computed by integrals of the form

$$\int \mathcal{D}\phi\, e^{iS_4(\phi)}\phi(x_1)\cdots\phi(x_n)$$

This expression is merely heuristic because the integral should be taken with respect to a measure $\mathcal{D}\phi$ that analogizes the Lebesgue measure to the space of classical trajectories for ϕ, and such a measure generally does not exist. Instead, physicists evaluate these integrals by generalizing the calculus of perturbed Gaussian integrals. Effective realists seek to interpret these perturbative calculations realistically by assimilating them to a realist account of quantization in general, such that the classical theory of the field ϕ is an approximately accurate description of the classical phenomena in the relevant domain.

These perturbative calculations naively lead to the divergent expressions that cause the realist's problems in Section 2. For example, when computing the amplitude $\langle\phi|\mathcal{M}|\phi\rangle$ for the scalar field to interact with itself you will encounter integrals of the form

$$\int d^4k \, \frac{1}{k^2 - m^2}$$

where k is a momentum. This integrand is the propagator for the scalar field ϕ, the inverse of the quadratic part of the action. The integral does not converge, seemingly giving a scattering operator that cannot be normalized. In practice, this problem is easily solved: replace the infinite integration limits with some energy Λ much larger than the characteristic scale m of this integral. This renders the integral finite, and the ensuing predictions are accurate.[2] But the facially *ad hoc* nature of this fix resists realist interpretation, as noted above.

Effective realists propose that we understand this truncation process realistically, as reflecting the theory's breakdown at high energy scales. By imposing an energy cutoff Λ we are neglecting high-energy contributions to the scattering operator, effectively replacing the space of all classical field trajectories with its subspace of trajectories with energy less than Λ. This is legitimated in the case of the scalar field because the theory's predictions are appropriately insensitive to the choice of Λ. That is, if ϵ is an energy scale smaller than Λ, then the scale-ϵ theory is obtained from the scale-Λ theory by removing trajectories with energy between ϵ and Λ and appropriately adjusting the measure.[3]

There is an obvious way to generalize the effective realist's proposal beyond this simple case, but this generalization leads to unacceptable consequences. For example, in QED the only new element is in accounting for spin. The relevant classical fields are a vector potential A_μ whose quanta are photons and a spinor field whose quanta are electrons and positrons. As in the scalar theory, scattering amplitudes in QED are computed by integrals of propagators; in the case of the photon field A_μ the propagator is

$$\frac{-\eta_{\mu\nu}}{k^2}$$

[2]Other sources of divergence remain in the theory, both at low energies and in divergent sums of infinitely many finite terms Miller (2021). In what follows I focus on the divergences associated with an ultraviolet regulator.

[3] Important for the effective realist's argument is the fact that the scale-ϵ action and the scale-Λ action have the same form (up to terms suppressed by Λ), differing only in the numerical values of masses and coupling constants. This result does not hold for the version of QED discussed in this section, but does hold for the theory of Section 7 (cf. Footnote 16).

with $\eta_{\mu\nu}$ the Minkowski metric.[4] And, as in the scalar theory, these integrals diverge and may be fixed by choosing an energy cutoff. Essentially the only difference is in the fact that A_μ is a massless vector field, making it transversely polarized. Its polarization is represented by a unit vector ϵ_μ living in a two-dimensional space of polarizations, and scattering amplitudes are linear in this polarization; for example, the amplitude for Compton scattering is of the form

$$\langle e^- \gamma | \mathcal{M} | e^- \gamma \rangle = \epsilon_\nu^* \epsilon_\mu \, \langle e^- \gamma^\nu | \mathcal{M} | e^- \gamma^\mu \rangle$$

where ϵ_μ is the polarization of the incoming photon and ϵ_ν the polarization of the outgoing photon. But the realist has no particular attitudes about polarization, so the rest of the effective realist story may proceed as before: we interpret this cutoff as a scale at which the theory breaks down, we restrict attention to the portion of the classical space of trajectories below that scale, and we interpret that literally.

If the effective realist proposal is generalized in this way, the resulting theory is nonunitary. Unitarity for the scattering operator constrains the imaginary part of the amplitude $\mathcal{M}$, and this in turn constrains the numerator of a propagator: it must be interchangeable with the sum over spin states when contracted with scattering amplitudes. (Schwartz, 2014, §24.1.3). This is a nontrivial requirement for the photon propagator, because its numerator $\eta_{\mu\nu}$ has full rank but the photon polarization vectors only span a two-dimensional subspace. Unitarity therefore requires the interaction amplitudes to remain in that two-dimensional subspace. Mathematically, this condition is expressed by the Ward–Takahashi identity (WTI):[5] for a photon with momentum p^μ, unitarity requires

$$p_\mu \, \langle \gamma^\mu \cdots | \mathcal{M} | \cdots \rangle = 0$$

Violations of the WTI can therefore serve as a measure of unitarity violations. And in cutoff QED, this identity is badly violated. For example, to leading order in the charge e and the cutoff Λ, the amplitude for a photon to interact with itself is

$$\langle \gamma^\mu | \mathcal{M} | \gamma^\nu \rangle = \frac{e^2}{2\pi^2} \Lambda^2 \, \eta^{\mu\nu}$$

On the effective realist interpretation, the cutoff Λ is effectively infinite. This quadratic divergence signals an extreme breakdown of unitarity.[6]

[4] This form of the propagator is responsible for the problem about to be raised, so it would deserve further justification if this section were really an argument against effective realism. However, any attempt at a derivation immediately encounters the interpretive issues to be discussed in the rest of this paper. Since the *reductio* of this section is designed to exhibit the practical consequences of these interpretive issues, readers who object to this propagator are likely already convinced of this section's conclusion.

[5] Various related results go by this name or the name "Slavnov–Taylor identity". For the purposes of the present argument, the relevant version of the WTI concerns the Lorentz structure of off-shell correlation functions.

[6] At this point expert readers are likely to start listing ways to reject this conclusion (an anonymous reviewer suggested at least six). Good: this is the point of this section. If you want to *modus tollens* this conclusion, then you agree that we have positive reason not to use a naive

I emphasize that this is a kind of nonunitarity that should worry the effective realist. Some kind of nonunitarity is to be expected from an effective theory, because effective theories ignore the probability to produce particles that only appear above the theory's scale of applicability. Unitarity should therefore be increasingly violated at high energies, where the neglected probabilities become appreciable. For example, the effective Euler–Heisenberg theory of light-by-light scattering at energies below the mass m of the electron gives the unpolarized cross section

$$\sigma(\gamma\gamma \to \gamma\gamma) = \frac{973e^8}{162000\pi^5}\frac{E^6}{m^8}$$

where m and e are the mass and charge of the electron, respectively, and E the scattering energy. As E increases this cross section grows quickly without bound, violating unitarity. But the breakdown of unitarity in cutoff electromagnetism is not of this kind. The violation of the WTI does not depend on the energy of the scattering interaction, so it gives unitarity violations at all energy scales. Moreover, this unitarity violation explicitly depends on the cutoff, which the light-by-light cross section does not, suggesting that the cutoff regulator is ultimately responsible for the non-unitarity. The realist cannot keep the worrying inconsistencies of QED behind a cutoff.

The empirically successful version of QED does not use cutoffs, and it is unitary, but it is not obviously amenable to the effective realist's proposal. In practice, physicists use a variety non-cutoff methods to evaluate the divergent integrals computing $\mathcal{M}$, preserving the WTI and hence unitarity (Bain, 2013; Georgi, 1993). One particularly convenient method is dimensional regularization, wherein divergent integrals are construed as meromorphic functions of the spacetime dimension d and divergences as their poles. When these function are combined with others to express an observable quantity as a function of spacetime dimension, this observable will have a removable singularity at $d = 4$, giving a finite prediction. So, like the cutoff method, dimensional regularization provides a well defined mathematical object at every step of the calculation. But unlike the cutoff method, dimensional regularization preserves the WTI, hence unitarity.

But dimensional regularization alone does not suffice to avoid the *reductio*. Effective realists will and should want to say that dimensional regularization (or Pauli–Villars, or whatever) is appropriate in this case and a simple cutoff is not. The question—in some ways the main question of this paper—is how they get to say it. Effective realism aims to give a prospective criterion of realist commitment, and for this it needs to say in advance when some regularization method is acceptable and when it is not (Fraser, 2018, 1166). But most presentations of effective realism deal only with cutoff regularization (Fraser 2018, 1170; 2020b,

---

ultraviolet cutoff in QED, and you can move on to the next section. The interesting question concerns the grounds on which it should be rejected. There are many natural suggestions, and they conflict. For instance, you may want to add a photon mass term to the Lagrangian to absorb this divergence (Peskin and Schroeder, 1995, 248), or you may want to disallow a cutoff because it violates gauge symmetry. These are inconsistent, because a mass term also violates gauge symmetry.

282; Ruetsche 2018, 1182; 2020, 297). When dimensional regularization is mentioned at all, the choice of regularization method is treated as a matter of convenience (Fraser, 2020a, fn. 10). And it's true that in simple theories like the scalar theory in Section 2, cutoff and dimensional regularization give the same results, and they both agree with many other regularization methods. Indeed, insensitivity to regularization methods is one criterion effective realists use to identify robust low-energy features of theories (Williams, 2019, 224). But this is exactly what fails for QED: different regularization schemes lead to different physics. QED is typical in this regard, so most QFTs resist effective realism.

Most presentations of effective realism depend on a realistic interpretation of a high-energy cutoff used to regulate perturbative QFTs. Taking these presentations too seriously leads to problems in most empirically relevant theories (namely, gauge theories), since these theories are nonunitary when cut off. The effective realist should not find this *reductio* compelling, because the obvious solution is to use a different regularization method. The purpose of the argument in this section is to show that sometimes an alternative regularization method is needed, and that its use cannot be justified on grounds of convenience. So if the effective realist wants to offer a prospective criterion of realist commitment, they need a story about why cutoff regularization should not be used in this case. And it's not obvious that they can do so, because specific features of the cutoff method apparently play an important role in the effective realist picture, both in delimiting the theory's domain of applicability and in characterizing sensitivity to the details of high-energy physics.

4 Realism and interpretation

The previous section argued that a particularly naive reading of the effective realist's proposal does not extend to empirically successful theories like QED. But perhaps a more careful reading does. In the rest of this paper I consider two possible ways of objecting to the *reductio* sketched in Section 3, both of which take the nonunitarity it encountered to be a symptom of a deeper inconsistency in cutoff QED. In both cases, this response is based on closer attention to the structure of the classical configuration space. These options are caricatures, meant to give an oversimplified illustration of how different semantic commitments can make a practical difference to the effective realist position. Indeed, both proposals are incomplete, and filling them in will likely require a more moderate semantic theory. In this section and the next, I argue that a conservative version of effective realism can adapt a common textbook argument to make the WTI a consequence of logical consistency. This identity is transmuted into a consistency condition by adopting a strong reading of the semantic tradition that underpins the selective realist strategy.

On a conservative reading of effective realism, it is a species of the currently popular "selective" strategy for scientific realism. We ought to believe in some but not all of a given theory's content, and our selection should be informed by the details of the theory under consideration (Fraser, 2020b; Ruetsche, 2020).

The great virtue of effective realism on this reading is the way that it unifies a positive argument for realism with an argument that some features of QFT will persist in future theories: the low-energy features of QFT are both responsible for its empirical successes and insensitive to the high-energy details that will be probed by future theories (Fraser, 2018, 2020b).[7]

Selective realism develops the realist's mid-century victories on the semantic flank of the realism–anti-realism debate. After an era of instrumentalist, operationalist, and conventionalist strategies for anti-realism, "[t]he current phase of the scientific realism debate… started in the middle 1960s and was based on an important consensus, viz., *semantic realism*" (Psillos, 2018, 20). On this consensus, there can be no dispute about what some scientific theory says: the meanings of theoretical statements are given by the same referential semantics that characterizes ordinary empirical descriptive discourse. Accepting the consensus doesn't mean accepting realism *simpliciter*, it just means turning one's anti-realist sentiment away from meanings and towards epistemic and metaphysical issues. On the selective realist consensus,

> [w]hen a scientist advances a new theory, the realist sees him as asserting the (truth of the) postulates. But the anti-realist sees him as displaying this theory, holding it up to view, as it were, and claiming certain virtues for it. (van Fraassen, 1987, 57)

What the scientist "holds up", in this metaphor, is the determinate semantic content of the theory: a collection of statements whose meaning is computed from the meanings of subsentential components by standard Tarskian methods. Terms for unobservable entities refer in the just same way that terms for observable entities do, and entities bear theoretical properties in just the same way that they bear non-theoretical ones. "Electrons have mass 0.5 MeV" is true just in case electrons have mass 0.5 MeV. The realist and anti-realist see these statements and agree to construe them all literally. They disagree about whether to believe only the statements about observable phenomena or something more.[8] Selective realists seek to characterize this something more, and conservative effective realism gives such a characterization: believe the literally construed statements that are insensitive to the value of the high-energy cutoff.

The debate over scientific realism therefore presupposes a supply of scientific statements that may be literally construed and then targeted by different attitudes. This is an idealization: scientific theories in the wild are not lists of claims,

[7]I will treat the versions of effective realism developed by J. D. Fraser (2018, 2020a,b) and Ruetsche (2018, 2020) as representative of the conservative strand of effective realism, and I will take Williams (2019) as an avatar of radical effective realism. This is an oversimplification. Fraser (2020b) expresses some hesitation about the standard account of interpretation, Williams (2019) consistently and explicitly places himself in the selective realist tradition, and Ruetsche is a sympathetic critic of effective realism who has tried to push it in a conservative direction. But I am interested in positions, not people. So I do not mean to disagree with Rivat's (2021, 12128) claim that Williams and Fraser can both be read as conservative.

[8]That is, on my reading, van Fraassen's reconciliation of semantic realism with empiricism is broadly analogous to Blackburn's (1993) and Gibbard's (2003) reconciliations of semantic realism with expressivism about moral talk.

they are networks of journal articles, textbooks, laboratory practices, scientists' attitudes, and more. That is, the statements to be literally construed in the realism debate are the product of interpretation. One of the great attractions of the selective realist consensus is that it allows this interpretation to precede realist debates. The work of interpretation can be done by realists or anti-realists and the results will be the same. Semantic realism puts no constraints on the task of interpretation, it only asks that the results admit a literal construal as truth-apt statements about the world. This issues in a picture of interpretation according to which

> to interpret a theory is to characterize the worlds possible according to it. These possible worlds are (i) models (in something like the logician's sense) of the theory, and (ii) characterized as physical (Ruetsche, 2011, 7).

Call this the "standard account" of interpretation. Some allowance must be made for the flexibility of scientific theories and the mathematics they use, and so to demand that the worlds of a theory be models in the strict logical sense would be asking too much. But this is no problem: all that selective realism requires from an interpretation is something enough like a logician's model that it admits a literal construal. And this is what the standard account of interpretation produces.

Implicit in this standard account of interpretation is the broadly logicist thesis that every mathematical object—or, at least, every mathematical object with a physical application—is sufficiently like a model in the logician's sense and that it represents its target systems through something more or less like model-theoretic satisfaction. This thesis fills the gap to which Ruetsche's parenthetical ambivalence is pointing. Strictly speaking, semantic realism is a thesis about the meaning of scientific statements. As such, it applies only to linguistically structured representations. But physicists do not formulate their theories directly in terms of statements about physical entities and their properties and relations. Instead, they use mathematical objects appropriate to their domain of investigation: Lorentzian manifolds, Hilbert spaces, and so on. But a commitment to standard referential semantics doesn't tell us anything about how to understand this. Literal construal can settle debates when physicists use words to say things about the world, but not when they use mathematics to say them. The logicist thesis converts the latter into the former by punning on the similarities between set membership and formula satisfaction, giving a sense in which mathematical representations may be given linguistic structure and hence literally interpreted.

As an illustration, consider a quantum particle in a one-dimensional box.[9] Possible states of this system may be represented by complex-valued functions on an interval like $[-1, 1]$. Since complex functions may be multiplied by complex numbers, any function ψ and any angle θ give another function $e^{i\theta}\psi$ satisfying $(e^{i\theta}\psi)(x) = e^{i\theta}\psi(x)$. If we have two particles in two boxes, one in state ψ and

[9] I am grateful to an anonymous reviewer for suggesting this example.

one in state $e^{i\theta}\psi$, then these two particles are in measurably different states. But if we consider only a single particle then ψ and $e^{i\theta}\psi$ name the same state. It is often said of this situation that while the "relative phase" of two systems is a real physical property, the "global phase" is not.

It is hard to make sense of this contrast between relative and global phases without assuming something like the logicist component of the selective realist consensus. In particular, it is hard to make sense of the claim that that "global phase" is not a real physical property, for there is nothing in the theory that the words "global phase" could be associated with. This is not the problem of negative existentials. Those are puzzling because they contain singular terms—which by nature purport to refer—whose successful reference the negative existential claim denies. So the truth of a negative existential seems to conflict with the contribution of its primary singular term. But there's nothing in quantum mechanics that even purports to refer to a "global phase".

But if you adopt the logicist component of the selective realist consensus, then talk of the global phase becomes meaningful—indeed, important. For there is a set of functions ψ that satisfy $\psi(0) = 1$. And by the logicist assumption, this set is the extension of some predicate P, which on a literal reading of the theory corresponds to some property. And if ψ names some state with this property, then $e^{i\theta}\psi$ must name a distinct state, since it does not belong to the extension of P. So on a literal reading, the states named by ψ and $e^{i\theta}\psi$ are distinct. And this makes the denial of the reality of the global phase an instance of selective realist commitment: read literally, the theory contains a property about which these two states disagree, but we are not committed to any physical differences between these states. By generalizing this example, the conservative effective realist obtains grounds on which to reject a naive cutoff in QED.

5 A pyrrhic solution

The conservative effective realist's semantic commitments give them grounds on which to say that Section 3's extension of effective realism to QED was incorrect, for it was based on a conceptual confusion. Recall that the breakdown of unitarity in Section 3 was signalled by a violation of the WTI. If the WTI merely expressed some feature that some QFTs happen to have, then its violation is problematic only insofar as it indicates a violation of unitarity. But there is reason to think that the WTI should itself be inviolable, for the WTI can be interpreted as expressing an exact symmetry. Exact symmetries often suggest "redundancy", because putatively distinct states they relate are equivalent for all practical purposes. But if the conservative effective realist would like to heed this suggestion, they may do so in only one way: by appealing to a theory in which the putatively distinct states are identified, thereby eliminating the distinction required by a literal reading of the original theory. And from the perspective of this new theory, the introduction of a cutoff in Section 3 is incoherent, for it rests on the distinction that will be eliminated in the "true" theory.

The textbook ban on cutoff regularization for electromagnetism and its generalizations is usually accompanied by something like the following justification (Peskin and Schroeder, 1995, §7.5). In classical electromagnetism, there are various reasons to think that the vector potential A_μ is a "redundant" description of the electromagnetic facts (Redhead, 2003; Healey, 2007). In particular, for any scalar function α the potentials A_μ and $A_\mu + \partial_\mu \alpha$ are essentially interchangeable: they have the same dynamical effect on charged classical and quantum particles, and if A_μ satisfies the laws of electromagnetism then so does $A_\mu + \partial_\mu \alpha$. This suggests that the potentials A_μ and $A_\mu + \partial_\mu \alpha$ in fact represent the same state of affairs. And if that's true, then it would be a straightforward contradiction for A_μ and $A_\mu + \partial_\mu \alpha$ to differ in any physically relevant way. But it is not hard to show that the interchangeability of A_μ and $A_\mu + \partial_\mu \alpha$ implies the WTI. This gives a justification for banning cutoff regularization: it violates the WTI, meaning that it distinguishes between A_μ and $A_\mu + \partial_\mu \alpha$, which represent the same state of affairs.

The conservative effective realist can adapt this argument to institute a similar ban on cutoff regularization for electromagnetism. For the effective realist, a QFT at energy scale Λ is defined by restricting to the space of trajectories with energy less than Λ. But if A_μ and $A_\mu + \partial_\mu \alpha$ represent the same point of the true electromagnetic configuration space, then there can be no subspace of electromagnetic configurations with momentum less than Λ, because we cannot make sense of electromagnetic configurations having a determinate momentum: by an appropriate choice of α, the potential $A_\mu + \partial_\mu \alpha$ can be made to have support at arbitrarily high momenta. Interpreted literally, cutoff regularization requires assigning inconsistent properties to the same physical configuration, so it should be rejected.

This strategy effectively rules out cutoff regularization, avoiding the *reductio* of Section 3. However, it rests on precisely the kind of promissory note that effective realism is meant to avoid. The idea—mooted, for example, by Guay (2008), Healey (2007), and Redhead (2003)—is that future mathematical developments will afford a re-description of the true electromagnetic configuration space that's acceptable to the conservative realist, one in which the state named by A_μ and $A_\mu + \partial_\mu \alpha$ is given a unique description. In this new theory, quantization will proceed as in other theories, using any of the usual methods for extracting finite results from field-theoretic data—any of the usual methods, that is, except a naive cutoff, which will not even be possible. And so in this new theory there will be no conflict between different regularization methods, evading the complaint of Section 3. But this new theory does not yet exist. And this leaves conservative effective realism in more or less the same place as older forms of realism, waiting in hope for the coming of some mathematical reconstruction of QFT with features conducive to realism.[10]

So conservative effective realism has the resources to block the argument of Section 3. If the future best interpretation of classical electrodynamics

[10] Worries about the promissory nature of this view are perhaps mitigated by the fact that there are other reasons to want new formulations of QFT—specifically, nonperturbative formulations. The strategy in this paragraph could be understood as a prediction about the nonperturbative treatment of gauge theories.

will replace A_μ and $A_\mu + \partial_\mu \alpha$ with a single mathematical representative of the common electromagnetic configuration they describe, then a literal interpretation of cutoff regularization attributes inconsistent properties to every electromagnetic configuration: it is on both the high- and low-energy sides of the cutoff. Because conservative effective realism partakes of the standard approach to interpretation, it takes literal interpretation as mandatory. So the argument of Section 3 is blocked on grounds of consistency. But the conservative effective realist is left without a positive realist story about QED.

6 Effective interpretation

Effective realists don't have to be conservative. Indeed, a more radical element in the effective realist camp rejects part of its selective realist heritage. In particular, radical effective realist arguments take aim at the logicist approach to characterizing the theoretical content of the mathematical component of a scientific theory that's part and parcel of the selective realist consensus (Williams, 2019). On this radical approach, the cutoff plays a role already in determining what the theory says, not just what sayings we should believe. Radical effective realism denies that interpretation is strictly prior to debates over realism, and it finds new philosophical resources in this denial. I will argue in the next section that these resources include a solution to the problem with cutoffs in electromagnetism.

Effective realists often argue for morals about interpretation in general. For example, Williams (2019, 211–212) objects to five principles apparently endorsed by most interpretive projects in the philosophy of physics:

(1) The theory to be interpreted is assumed to provide a true and exhaustive description of the physical world in all respects, including at all length scales.

(2) A theory is to be interpreted in isolation. . . .

(3) An interpretation of a theory consists of the set of all worlds nomologically possible according to that theory.

(4) This set of possible worlds is determined by generic structural features of the theory in question. . . . Information about empirical applications of the theory. . . [is] largely or entirely ignored.

(5) The goal of interpreting a physical theory is to identify and characterize its fundamental ontological features.

As Williams argues, adhering to these principles would make the interpretation of effective QFTs difficult. The intended scale dependence of these theories sits uncomfortably with principles (1) and (2), while the particular mathematical form of the quantum theories found in practice resists principles (3) through (5). As a result, philosophical interpreters of high-energy physics look to alternative, in-development formalisms that are susceptible to the above principles, instead of the standard theory as applied in the laboratory. And this untethers realism from the empirical successes that motivate it—a fate also met by conservative effective realism in Section 5.

The principles above are explicable—perhaps even compelling—if we see them as features of the selective realist consensus. Consider the first two principles, according to which a theory is taken to provide a true and self-contained description of the world for the purposes of interpretation. On the selective realist consensus, a theory would be true if the world were just exactly as the theory says it is. And on this consensus, "what the theory says" is something about which there can be no debate, for we have all agreed to compute the meanings of scientific statements from their subsentential components using standard referential semantics. The results of interpretation are therefore something that may be "held up to view", in van Fraassen's idiom: a fixed body of claims with determinate semantic content toward which realists and anti-realists can take different attitudes of belief or acceptance. An incomplete interpretation fails to completely separate the interpretive question of what the theory says from the metaphysical question of whether the theory is true, and a non-literal interpretation is backsliding into the bad old days of operationalism and *Protokollsätze*.[11]

So on a radical reading, we can take Williams's criticisms of the above principles to motivate a rejection of some part of the selective realist consensus. For example, consider Williams's disagreement with D. Fraser (2009) over the interpretation of lattice regularization, in which divergences are tamed by representing space as a lattice of spacetime points, each a fixed distance L from its neighbors. Fraser (2009, 552) argues that a theory using lattice regularization is not a serious candidate for a realistic interpretation of QFT. On a literal interpretation of lattice QFT, there is a fundamental minimum length L, making spacetime discrete. And this lattice structure seems necessary for finite predictions. So realism about this theory involves the belief that spacetime is discrete. Since no one believes this to be an implication of QFT, lattice QFT is not a serious candidate for realism. Assuming the selective realist consensus, this argument is impeccable. Williams (2019, 217) recasts Fraser's argument as a *reductio* in which the standard account of interpretation is an explicit premise and urges us to reject it, instead of lattice QFT.

That radical effective realism specifically rejects the semantic realist's logicist account of mathematical representation—rather than, say, the broadly Tarskian semantics for theoretical claims—is illustrated by Williams's (2019, 222) analysis of this case. From his response to Fraser's argument, it's at least clear that he does not take lattice QFTs to impute a discrete structure to spacetime; this alone is a flat rejection of the standard account's logicist thesis. But he also gives a positive argument against a logicist interpretation of other features of the theory. Roughly speaking, the Nielsen–Ninomiya theorem says that applying

[11] Principles (3)–(5) also follow from semantic realism about specifically modal language in terms of referentialist—that is, possible worlds—semantics. This list also illustrates the range of possible semantic commitments that might lie between the caricatured conservative and radical poles I am discussing: one might in principle hold principles (1) and (2) but not (3)–(5), or some other combination.

lattice regularization to a theory in four dimensions results in a regulated theory with sixteen times as many types of fermions as the unregulated theory. Read according to the semantic realist's policies, the regulated theory says that lattice QFT is incapable of representing the world as containing just an electron, without fifteen "mirror fermion" partners. As Williams (2019, 223) points out, the mathematical features of these mirror fermions depend on the length cutoff L. As a result, he argues, we can and should understand this theory as representing a world containing only one type of fermion, not sixteen.

It emerges from this contrast that the energy cutoff introduced in Section 2 plays a different role for the conservative and radical effective realists.[12] For a conservative effective realist, the cutoff offers a criterion of selection that may be applied to an interpreted theory. The theoretical content of the theory is completely fixed by semantic realism and a linguistic treatment of mathematical objects, and the cutoff guides our selection from this menu of theoretical contents. In the case of lattice QFT, for example, it seems that the conservative effective realist is committed to saying that fermions really must come in sexdecuples. Some properties of these mirror fermions may not deserve realist commitment, given their cutoff dependence, but their existence is a cutoff-independent fact. Indeed, Ruetsche (2020, 311) presses exactly this problem. But it is only a problem for the conservative, with their commitment to the standard account of interpretation. For the radical effective realist, the cutoff is a tool for justifying an identification of the theoretical contents of the theory in the first place.

7 Electromagnetism, effectively interpreted

Radical effective realism affords a unitary and scale-relative interpretation of QED. Like the conservative, the radical effective realist sees an important difference between the photon field A_μ and ordinary scalar fields, and so the *reductio* of Section 3 incorrectly generalized the method of cutoff regularization to the configuration space of the photon field. The problems of the regularization of Section 3 came about because the cutoff was incorrectly implemented, and so it did not actually separate low- and high-energy modes. At least, this is what the radical effective realist can say. They can say this because they reject the logicist component of the selective realist consensus and its restrictions on the use of mathematics. In place of logicism, the radical effective realist offers case-specific analyses of specific uses of mathematics, resulting in interpretations that might not be neutral on the representational role of the mathematics' parts. This permits a version of cutoff electromagnetism that's unavailable to the selective realist consensus.

The radical effective realist faces two problems. They need a response to the *reductio* of Section 3. And in light of Section 5, they need to account for the claim that A_μ and $A_\mu + \partial_\mu \alpha$ represent the same state of affairs, at least in some sense. They must reject the story of Section 5 if they are to save the scale-relativity of electromagnetism. But the claim that A_μ contains "redundancy" plays a

[12] This difference is also emphasized by Rivat (2021, 12127).

broad role in QED, and a scale-relative understanding of electromagnetism must be compatible with this role. Fortunately for the radical effective realist, their conception of compatibility is significantly more flexible than the conservative's. It only demands case-by-case accounts of mathematics in a physical theory, and these accounts are given by the mathematics' use, rather than by a linguistic interpretation of its structured-set realization.

So to make sense of redundancy talk, the radical effective realist must look to its use in QED. The idea that A_μ is a "redundant" representation was introduced in Section 5 to argue for the conceptual necessity of the WTI, but this identity is downstream from other uses. The potentials A_μ and $A_\mu + \partial_\mu \alpha$ give two-point functions that differ by a term

$$p_\mu p_\nu \langle \gamma^\mu | \mathcal{M} | \gamma^\nu \rangle$$

If these potentials represent the same state of affairs, then they must agree on physically significant quantities like the two-point function. And so the difference term must vanish, as the WTI says. But classical potentials can only "give" a two-point function if they are set within some framework for deriving quantum predictions from classical descriptions. And appeals to redundancy are already used in setting up this framework.

The precise role of redundancy talk in the quantization of electrodynamics is disputed, but generally speaking it is used to solve a problem with applying standard quantization techniques to theories like classical electromagnetism (Redhead, 2003; Healey, 2007; Dougherty, 2021). Heuristically, the path integral is dominated by stationary points of the action, and so to lowest order will reproduce the predictions of a classical field theory with the same action and configuration space. Maxwell electromagnetism is therefore a natural guide in choosing a configuration space and action for QED. However, the usual Maxwell action

$$S_M(A) = \int d^4x \, \frac{1}{2} A_\nu (\eta^{\mu\nu} \partial^2 - \partial^\mu \partial^\nu) A_\mu$$

gives an ill-defined path integral: every amplitude diverges, even if you introduce a regulator. It is not hard to show that this further divergence is due to the symmetry of this action under the transformation sending A_μ to $A_\mu + \partial_\mu \alpha$. Informally, it is said that the amplitude "is badly defined because we are redundantly integrating over a continuous infinity of physically equivalent field configurations" (Peskin and Schroeder, 1995, 295).

The use of this redundancy diagnosis is in suggesting a treatment. However, the connection between redundancy talk and practical computations is not straightforward. In particular, the conservative effective realist cannot understand the practical response to this divergence problem as an elimination of redundancy. On the selective realist consensus, there's only one thing redundancy could mean: the potentials A_μ and $A_\mu + \partial_\mu \alpha$ are discrepant descriptions of the same state of affairs, and so they impute contradictory properties to a single physical configuration. A consistent, non-redundant formulation of the classical theory will be a new theory in which there will be some single description that includes only properties on which A_μ and $A_\mu + \partial_\mu \alpha$ agree. And indeed, the conservative effective realist appealed to a promised reformulation of just this kind in Section 5 when banning the use of cutoff regularization in electromagnetism. But this is not how the divergence is dealt with.

In practice, "eliminating the redundancy" means explicitly equipping the classical configuration space with mathematical structure that I will call "gauge structure". The function of gauge structure is to represent the fact that transformations like $A_\mu \mapsto A_\mu + \partial_\mu \alpha$ are not mere symmetries of the Maxwell action. That is, invariance under such transformations is not a feature had by some functions on the electromagnetic configuration space and not others and which the Maxwell action happens to exhibit. Rather, appropriate invariance—expressed by compatibility with the gauge structure—is a requirement for a coordinate expression on the configuration space to define a function, in the same way that it must give a unique output for any input. Mathematically, this requirement is implemented by a configuration space that includes coordinates parametrizing the distinguished transformations. A point in this configuration space is coordinatized by a pair (A_μ, c), where the vector field A_μ parametrizes configurations of the electromagnetic field and the scalar skew-commuting field c parametrizes the space of distinguished transformations $A_\mu \mapsto A_\mu + \partial_\mu \alpha$. This means that the configuration space is equipped with an operator δ acting as

$$\delta A_\mu = \partial_\mu c \qquad\qquad \delta c = 0$$

reflecting the action of the transformations on A_μ and the abelian nature of the group of transformations.[13] For some expression S in A_μ and c to give a well defined function on the configuration space, it must satisfy $\delta S = 0$. In particular, any classical action must satisfy this requirement if it is to be used in the path integral.[14] The δ operator also determines identity criteria for functions on this space: two expressions S and S' name the same function if $S - S' = \delta h$ for some expression h.

One practical role of gauge structure is to permit alternative expressions for the classical action. The Maxwell action S_M gives a singular path integral on the (A_μ, c) configuration space because it does not depend on the c coordinate. This is a coordinate singularity that can be resolved by using a different expression for the same function on configuration space—that is, by choosing some h such that $S_M + \delta h$ does depend on c. The simplest example is the expression

$$S_F(A, c) = \int d^4x \left[\frac{1}{2} A^\mu \partial^2 A_\mu - \bar{c}\, \partial^2 c \right]$$

Once this change of coordinates has been effected, this theory can be quantized just as you would quantize a theory without gauge structure, using regulated

[13] In the nonabelian case, the operator δ acts as $\delta A^a_\mu = \partial_\mu c^a + f^a{}_{bc} A^b_\mu c^c$ and $\delta c^a = -\frac{1}{2} g f^a{}_{bc} c^b c^c$, with g the coupling constant of the theory. See Peskin and Schroeder (1995, 16.4) for a more detailed discussion.

[14] The path integral measures $\mathcal{D}A_\mu$ and $\mathcal{D}c$ must also be appropriately compatible with δ. The mutual compatibility of δ with the action S and measures $\mathcal{D}A_\mu$ and $\mathcal{D}c$ is captured by the Batalin–Vilkovisky quantum master equation (QME) (Costello, 2011, §5.1.2). In particular, the WTI follows from the QME.

integrals as in Section 2 or more general frameworks (Costello, 2011, Ch. 2). In particular, this theory can be regulated using a sharp cutoff without encountering any problematic unitarity. Imposing a cutoff is a coordinate-dependent procedure, and Section 2 did it in the wrong coordinates.

Indeed, gauge structure furnishes the radical effective realist with solutions to both of their problems. To account for the claim that A_μ and $A_\mu + \partial_\mu \alpha$ represent the same state of affairs, they need only point to the operator δ and its use in replacing the action S_M with S_F. This structure is the cash value of "redundancy" talk, and according to the radical effective realist we needn't say anything more. Of course, there might be further specific interpretive questions about the relationship of redundancy talk to the WTI or to cutoff renormalization. But when it comes to redundancy talk itself, the radical has a complete account—an account the conservative can't share, because it says that A_μ and $A_\mu + \partial_\mu \alpha$ represent the same state of affairs without reformulating the theory as semantic realism demands.

More importantly, the explicit representation of the equivalence of A_μ and $A_\mu + \partial_\mu \alpha$ allows for its explicit scale-relativity. The equations introduced in the paragraph before last—such as the relationship $\delta A_\mu = \partial_\mu c$ between the fields and the requirement $\delta S = 0$ on the action—make demands at every scale. But they don't make demands across scales. That is, the equation $\delta S = 0$ is nontrivial for arbitrarily high energies, but it does not enforce correlations between high and low energies. It follows that we can relativize the operator δ to an energy scale Λ, giving an operator δ_Λ. Relativizing the condition on the action gives the requirement $\delta_\Lambda S_\Lambda = 0$ on the action at scale Λ, expressing the scale-Λ equivalence of A_μ and $A_\mu + \partial_\mu \alpha$.[15] The scale relativity of effective realism can thereby be reconciled with the redundancy talk used in quantization, the justification of the WTI, and so on.

The radical effective realist can use the scale-relativity of gauge structure to resolve the problems with cutoff electromagnetism that I raised in Section 3. To properly cut out the high energy regions electromagnetic configuration space, it's not enough to simply discard features of the electromagnetic potential above some energy cutoff. The gauge structure on the configuration space must also be relativized, which leads to a scale-relative WTI—that is, a version of the WTI that permits unitarity violations only beyond the scale of the theory. This scale-relative identity is enough to save scale-relative unitarity, solving the problem of Section 3. The conceptual problem of high-energy sensitivity is also resolved. The renormalization group equation ensures that the low-energy predictions of the theory are robustly independent of the details of the high-energy effective actions. Crucially, the condition $\delta_\Lambda S_\Lambda = 0$ holds at one energy scale Λ if and only if it holds at every scale, so the radical effective realist can account for the redundancy talk surrounding gauge structure without committing to any high-energy details that don't follow from low-energy ones (Costello, 2011, Lemma 9.2.2).[16]

[15]More carefully, what is relativized to scale Λ is the QME. Since the QME bundles in a generalization of the WTI, this gives a scale-relative version of the identity, as well.

[16] In particular, this fact permits the effective realist's coarse-graining story, because gauge

Because the radical effective realist is not bound by any particular thesis about how physicists use mathematics to represent the physical world, they can (and must) have case-specific stories about how to coordinate the mathematical manipulations with theoretical content. In cases where this story agrees with the conservative's, there's rarely a need to be explicit about it. But in cases like QED these stories diverge, and the actual mathematical incarnation of redundancy talk is only available to the radical. Once the mathematical structure representing this redundancy talk is identified, it can be relativized to energy scales and the effective realist's scale-relative story goes through.

8 Conclusion

On a naive reading, effective realism has a problem with empirically successful theories. Realists seek to explain empirical success by reference to descriptive accuracy, and effective realism is meant to enable such explanations by appealing to effective field theory methods. These methods get their power from their ability to isolate physically relevant degrees of freedom at a given scale and to identify degrees of sensitivity to the details of physics at higher energies. Effective realists want to harness this power to create a version of scientific realism that's applicable to our best physical theories, explanatory of these theories' successes, and robust against future theory development. However, the most direct extension of effective realism to the most successful theories is self-undermining: restricting to any energy scale results in a non-unitary theory, which by the effective realist's lights means we have left the domain of empirical applicability.

The possible responses to this conflict illustrate a division within the effective realist program. A conservative response appeals to the selective realist consensus and the standard account of theory interpretation to make the WTI trivially true: a consequence of the redundancy of a representation in which both A_μ and $A_\mu + \partial_\mu \alpha$ appear. But this leaves the effective realist with no realist story about QED. A more radical response takes up some of the effective realists' criticisms of the standard account of theory interpretation, rejecting the conservative's interpretation of redundancy talk and looking to its use in computation. The radical's more flexible semantic commitments allow them to interpret redundancy talk in a way that allows for scale-relativity, and from this a scale-relative WTI follows.

One proximal conclusion of this discussion is that the conservative and radical tendencies of effective realism are different, and this difference makes a difference. Another is that the radical is on better ground, at least with respect to treating electromagnetism effectively. It's less clear that this radicalism can sustainably be extended to a full scientific realism, or that effective realism will be as attractive if purged of its conservative elements. After all, the selective realist consensus is a consensus for a reason. But as the case of redundancy talk shows, the

structure is preserved by coarse-graining in the same way that the form of the classical action is. So the result of Footnote 3 generalizes (Costello, 2011, Ch. 6).

radical effective realist need not object to the standard account of interpretation wholesale. The radical effective realist can be a committed semantic realist when it comes to the propositional content of the theory: unobservables like pions exist, they have theoretical properties like spin and charge, and these unobservables and their properties account for the patterns we find in the phenomena. The disagreement only concerns the extraction of propositional content from the theory's mathematics. And here it must be admitted that the standard account's appeal to "something like a logician's model" is, for all its success, something of a stopgap measure. We should join the radical effective realist in rejecting it.

References

Bain, J. (2013). Effective field theories. In Batterman, R., editor, *The Oxford Handbook of Philosophy of Physics*, pages 224–254. Oxford University Press.

Blackburn, S. (1993). *Essays in Quasi-Realism.* Oxford University Press.

Boyd, R. (1989). What realism implies and what it does not. *Dialectica*, 43(1-2):5–29.

Chakravartty, A. (2007). *A Metaphysics for Scientific Realism.* Cambridge University Press.

Costello, K. (2011). *Renormalization and Effective Field Theory.* American Mathematical Society.

Dougherty, J. (2021). I ain't afraid of no ghost. *Studies in History and Philosophy of Science*, 88:70–84.

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

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

Fraser, J. D. (2018). Renormalization and the formulation of scientific realism. *Philosophy of Science*, 85(5):1164–1175.

Fraser, J. D. (2020a). The real problem with perturbative quantum field theory. *The British Journal for the Philosophy of Science*, 71(2):391–413.

Fraser, J. D. (2020b). Towards a realist view of quantum field theory. In French and Saatsi (2020), pages 276–292.

French, S. and Saatsi, J., editors (2020). *Scientific Realism and the Quantum.* Oxford University Press.

Georgi, H. (1993). Effective field theory. *Annual review of nuclear and particle science*, 43(1):209–252.

Gibbard, A. (2003). *Thinking How to Live.* Harvard University Press.

Guay, A. (2008). A partial elucidation of the gauge principle. *Studies in History and Philosophy of Modern Physics*, 39(2):346–363.

Halvorson, H. and Müger, M. (2007). Algebraic quantum field theory. In Butterfield, J. and Earman, J., editors, *Philosophy of Physics Part A*, pages 731–922. Elsevier.

Hancox-Li, L. (2015). Coarse-graining as a route to microscopic physics: the renormalization group in quantum field theory. *Philosophy of Science*, 82(5):1121–1123.

Hanneke, D., Fogwell, S., and Gabrielse, G. (2008). New measurement of the electron magnetic moment and the fine structure constant. *Phys. Rev. Lett.*, 100:120801.

Healey, R. (2007). *Gauging What's Real.* Oxford University Press.

Hesse, M. (1976). Truth and the growth of scientific knowledge. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 2, pages 261–280.

Kuhlmann, M. (2010). Why conceptual rigour matters to philosophy: On the ontological significance of algebraic quantum field theory. *Foundations of Physics*, 40(9):1625–1637.

Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science*, 48(1):19–49.

Miller, M. E. (2021). Infrared cancellation and measurement. *Philosophy of Science*, 88(5):1125–1136.

Peskin, M. E. and Schroeder, D. V. (1995). *An Introduction to Quantum Field Theory.* Addison–Wesley.

Polchinski, J. (1984). Renormalization and effective Lagrangians. *Nuclear Physics B*, 231(2):269–295.

Psillos, S. (2005). *Scientific Realism.* Routledge.

Psillos, S. (2018). The realist turn in the philosophy of science. In Saatsi, J., editor, *The Routledge Handbook of Scientific Realism*, pages 20–34. Routledge.

Putnam, H. (1975). What is mathematical truth? In *Mathematics, Matter and Method.* Cambridge University Press.

Redhead, M. (2003). The interpretation of gauge symmetry. In Brading, K. and Castellani, E., editors, *Symmetries in Physics*, pages 124–139. Cambridge University Press.

Rivat, S. (2019). Renormalization scrutinized. *Studies in History and Philosophy of Modern Physics*, 68:23–39.

Rivat, S. (2021). Effective theories and infinite idealizations: a challenge for scientific realism. *Synthese*, 198(12):12107–12136.

Rueger, A. (1990). Independence from future theories: A research strategy in quantum theory. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pages 203–211. Philosophy of Science Association.

Ruetsche, L. (2011). *Interpreting Quantum Theories*. Oxford University Press.

Ruetsche, L. (2018). Renormalization group realism: The ascent of pessimism. *Philosophy of Science*, 85(5):1176–1189.

Ruetsche, L. (2020). Perturbing realism. In French and Saatsi (2020), pages 293–314.

Schwartz, M. D. (2014). *Quantum Field Theory and the Standard Model*. Cambridge University Press.

Stanford, P. K. (2006). *Exceeding Our Grasp*. Oxford University Press.

van Fraassen, B. C. (1987). *The Scientific Image*. Clarendon Press.

Williams, P. (2015). Naturalness, the autonomy of scales, and the 125 GeV Higgs. *Studies in History and Philosophy of Modern Physics*, 51:82–96.

Williams, P. (2019). Scientific realism made effective. *The British Journal for the Philosophy of Science*, 70(1):209–237.

Williams, P. (2021). Renormalization group methods. In Knox, E. and Wilson, A., editors, *The Routledge Companion to Philosophy of Physics*, chapter 19. Routledge.

Wilson, K. G. and Kogut, J. (1974). The renormalization group and the ϵ expansion. *Physics reports*, 12(2):75–199.

Worrall, J. (1989). Structural realism: The best of both worlds? *Dialectica*, 43(1-2):99–124.