ARTICLE

Λ and the Limits of Effective Field Theory

Adam Koberinski[1,2]* and Chris Smeenk[3]

[1]Department of Philosophy, University of Waterloo, Waterloo, ON, Canada, [2]Center for Philosophy of Science, University of Pittsburgh, Pittsburgh, PA, USA and [3]Department of Philosophy and Rotman Institute of Philosophy, University of Western Ontario, London, ON, Canada
*Corresponding author. Email: adam.koberinski@pitt.edu

Abstract

The cosmological constant problem stems from treating quantum field theory and general relativity as an effective field theory (EFT). We argue that the problem is a reduction ad absurdum and that one should reject the assumption that general relativity can generically be treated as an EFT. This marks a failure of naturalness and provides an *internal* signal that EFT methods do not apply in all spacetime domains. We then take an *external* view, showing that the assumptions for using EFTs are violated in general relativistic domains where Λ is relevant. We highlight some ways forward that do not depend on naturalness.

1. Introduction

Quantum field theory (QFT) provides a general framework for formulating physical theories, replacing predecessors with similar scope, such as classical Lagrangian mechanics. Physicists have developed successful QFTs for the weak, strong, and electromagnetic forces, but what are the prospects for gravity? Early efforts to formulate a QFT for gravity showed that it lacked a feature then taken as necessary for a sensible QFT: perturbative renormalizability. For theories with this property, such as QED, the infinities that arise in calculating quantities through a perturbative expansion around the free field theory can be tamed by reparametrizing a finite number of "bare" coupling parameters appearing in the Lagrangian. The renormalized theory then yields predictions regarding diverse physical processes. A perturbative expansion of general relativity (GR) differs strikingly from QED (and theories of the other forces), however, because of the dimension of its coupling constant. Heuristic "power-counting" arguments link the dimension of the coupling constant(s) to the ultraviolet behavior of the theory and suggest that no finite reparametrization will eliminate all of GR's ultraviolet infinities. These arguments have been supplemented by rigorous proofs that gravity fails to be perturbatively renormalizable.[1]

---

[1] Goroff and Sagnotti (1986) prove that a perturbative expansion of vacuum GR around Minkowski spacetime diverges at two loops; divergences arise already at one loop for gravity coupled to matter

Yet these results no longer present a roadblock, given the dramatic reversal of fortune non-renormalizable theories have experienced. This new perspective follows from the use of renormalization group techniques to clarify how different terms in a Lagrangian behave under changes of scale. Predictions can still be extracted from (some) non-renormalizable Lagrangians, whose low-energy properties can be fully characterized in terms of a finite set of parameters.[2] Physicists routinely construct effective field theories (EFTs) designed to mimic the low-energy physics of more fundamental Lagrangians. The finite set of parameters sufficient to specify low-energy behavior (e.g., coupling constants and masses) can then be determined experimentally, leading to a variety of further predictions, just as in the case of renormalizable theories. The presence of non-renormalizable term(s) in the Lagrangian, rather than indicating a failure, merely delimits the domain of applicability of the EFT. Furthermore, many distinct candidates for a (more) fundamental Lagrangian may generate the same low-energy EFT. When this is the case, the EFT is insensitive to contrasts in the descriptions of higher-energy physics they provide. As Ruetsche (2020) succinctly puts it, "*T* is merely effective just in case *T*, while not itself a complete and accurate account of physical reality, approximates that account *whatever it is (!)* within a restricted domain of application" (298, original emphasis).

The EFT approach promises to justify our confidence in low-energy theories while remaining agnostic about physics at higher-energy scales. Making good on this promise requires some assurance that all reasonable candidates for a (more) fundamental theory flow to a low-energy EFT. The cases where physicists have been able to prove that the renormalization group flow has the desired properties share two features: locality and naturalness. Locality is the requirement that the Lagrangian depends on fields and their derivatives at a point. Although naturalness has been used in a variety of distinct senses, Williams (2015) argues convincingly that these can all be seen as stemming from the concept of autonomy of scales: the expectation that physics at low-energy scales decouples from physics at higher energies. If naturalness holds, the dynamics within the relevant domain are insensitive to the details of physics at higher-energy scales. Although often left unstated, there are some minimal structures required to set up an EFT, such as a method for demarcating high- from low-energy degrees of freedom. At a minimum, this requires enough spacetime structure to define a useful notion of energy and a sufficiently strict division between high and low energies, the latter of which falls within the domain of the EFT. We will discuss these issues further (section 3), but these brief comments are sufficient to illustrate that the criteria for a (more) fundamental theory to be well approximated by a low-energy EFT are much less restrictive than those imposed by demanding a renormalizable QFT. This suggests a very different take on the "problem of quantum gravity": To what extent can we treat classical general relativity as the low-energy EFT of an unknown quantum theory?

---

fields. However, the asymptotic safety program is committed to the idea that a nonperturbative renormalization of GR leads to a well-defined theory for all energy scales. See Friederich (2018) for an introduction aimed at philosophers.

  [2] These are theories whose high-energy Lagrangians flow under the action of the renormalization group to a finite-dimensional subspace in the space of possible theories, characterized by a finite collection of parameters. See, for example, section 16.4 in Duncan (2012) for a technical discussion and Ruetsche (2020) for a philosophical overview.

Indeed, EFT methods have been successfully applied to a variety of problems in gravitational physics over the last two decades (see, e.g., Burgess 2004; Donoghue 2012, for reviews). However, a careful analysis of the domains in which EFT methods work for gravity highlights the exceptional nature of these cases. From successful applications, we learn that EFT methods work for models that can be treated as "nearly" static or (asymptotically) flat, but they do not work in a variety of other situations routinely described with classical gravity. In attempting to construct an EFT for dynamically evolving models in cosmology, for example, self-consistency problems arise in assessing whether one has actually expanded around a solution (see, in particular, Bianchi and Rovelli (2010) and further discussion in section 3). The more general question of whether *all* gravitational models can be treated using EFT methods remains open.

One response takes the applicability of EFT methods as a new criterion of adequacy: if we cannot construct an EFT, then we have no way of understanding how to treat a classical solution as an approximation to a more fundamental quantum model. Although looking for keys lost at night under the lamppost is often a good strategy, this response seems to foreclose the possibility of a further generalization, like the move from renormalizable QFTs to EFTs. The need for further generalization would not be a surprise: research programs in quantum gravity have had to replace the spacetime structures employed in formulating conventional QFT—such as Poincaré symmetries and the causal structure of Minkowski spacetime—with structures definable in generic curved spacetimes. Here, we will assess the assumptions of the EFT framework and argue that they also impose constraints that gravity might force us to break. The focal point for our discussion is the cosmological constant problem (CCP), which we take to signal the internal breakdown of EFT methods for gravity, particularly over cosmic distance scales in near-Friedmann–Lemaître–Robertson–Walker (FLRW) spacetimes.

Suppose that (i) we treat classical GR as the lowest-order term in an EFT, whose action S^{eff} in principle follows from a full theory of quantum gravity via integrating out higher-energy modes from the "true" action S. We assume that the Planck mass is the energy scale used to separate high- from low-energy degrees of freedom, with S^{eff} only concerned with the latter. If we assume that (ii) this EFT is stable and autonomous with respect to higher-energy physics and is able to reproduce *all* effects of classical GR, trouble arises as a result of the relevant terms in the effective action:[3]

$$S^{\text{eff}} = \int \sqrt{-g}\, d^4x \left(-\Lambda + \frac{m_P^2}{2}R + c_1 R^2 + c_2 R_{\mu\nu}R^{\mu\nu} + \cdots \mathcal{L}_m \right) \tag{1}$$

$$= \int \sqrt{-g}\, d^4x \left(-\Lambda + \frac{m_P^2}{2}R + \sum_{n=0}^{\infty} \sum_{i} \frac{c_i}{m_P^{2n}} \mathcal{O}_i^{[2n+4]} + \mathcal{L}_m \right). \tag{2}$$

We have made explicit the nature of the EFT expansion for the gravitational terms in equation (2). The first two terms are the familiar Einstein–Hilbert action terms. The $\mathcal{O}_i^{[2n+4]}$ are higher-order terms in the gravitational Lagrangian with mass dimension

---

[3] In this expression, R is the Ricci scalar; Λ is the cosmological constant; $\mathcal{L}_m$ is the Lagrangian for any matter fields; $m_P = 1/(8\pi G)$ is the Planck mass in units, where $c = \hbar = 1$; and g is the determinant of the metric g_{ab}.

$2n + 4$, constructed from the Riemann and Ricci tensors and Ricci scalar and subject to the symmetry constraints of GR. The terms are ordered by their mass-energy dimension; constants c_i are dimensionless coupling strengths, suppressed by the explicit powers of the Planck mass. At low energies relative to the Planck mass, higher-dimension terms will be heavily suppressed. Note that the cosmological constant term Λ has mass dimension 4. A similar expansion in the matter Lagrangian leads to a full EFT treatment of gravity and matter.

If we assume (iii) that the couplings in this EFT vary under renormalization group flow, we run into a problem. Both the Λ term and the first term in the matter EFT expansion $\langle \rho \rangle$ have mass dimension 4 and are relevant parameters; they fail to be "natural" in that they receive contributions proportional to m^4 under renormalization group flow from one energy scale m down to another. Even if we stipulate that Λ has a small value in an effective action S' at some high-energy scale, this will not be true for the action S^{eff} obtained at a lower scale via the renormalization group as a result of to radiative corrections.[4] When we write the EFT action, we also assume (iv) that the zero-point energies minimally couple to gravity. The zero-point energies from the quantum fields then take the same constant form in the action as Λ, so these terms should be grouped together. If this EFT applies everywhere below the cutoff scale, then this $\Lambda + \langle \rho \rangle$ term would have various observable effects (described later in the article). Given the quartic dependence of both on the mass-energy scale m, even integrating from relatively low-energy scales leads to a dramatic conflict—"the worst prediction in the history of physics"—with observational bounds.

This is, in a nutshell, the CCP.[5] What should we make of it? Here, we want to draw a contrast between three different responses and, in particular, explore the third:

1. *Anthropic parameter fixing*: Accept all previously noted assumptions except for (iii), and reconsider how to think of parameters like Λ (along with other apparently finely tuned aspects of low-energy physics). Specifically, we should take the observed values as "anthropically selected" from an ensemble of possible values. (This strikes us as an act of desperation.)

2. *Modify dynamics; keep EFT*: Accept assumptions (i)–(iii), and reject (iv). Although the EFT concepts apply, we have overlooked something that will change the problematic scaling behavior (e.g., supersymmetry, change in the number of dimensions of spacetime, modified gravity, etc.). Thus, we should modify the particular details of gravitational dynamics so that the EFT framework applies.[6]

---

[4] Although it is possible to arrange a delicate cancellation between the "bare" value of the parameter and these radiative corrections, that is unappealing. Physicists often refer to such cancellations derisively as "fine-tuning." In addition, the radiative instability of the cosmological constant term means that the Λ must be fine-tuned again and again at higher orders (cf. Koberinski (2021a)).

[5] See Martin (2012) for a thorough review of the CCP, and see Rugh and Zinkernagel (2002) for a philosophical discussion.

[6] We note that some solutions fitting into this camp slightly modify assumptions (i) or (ii). For example, $f(R)$ theories modify (i) because there are further classical terms that should be included in the zeroth-order EFT. The main point here is that these approaches largely accept the full applicability of the EFT approach.

3. *Reject EFT*: The argument is a reductio ad absurdum of the combination of assumptions (i) and (ii), namely, that we can treat all of classical GR as a low-energy EFT. The resulting failure of naturalness signals an internal inconsistency with the application of EFT methods to some specific domains of gravitational physics.

In pursuing the third line, we step back to take a look at the assumptions required to set up an EFT for gravity.[7] We find that for spacetimes where one would expect Λ to have a significant effect, we cannot set up a well-defined separation of energy scales. The failure of naturalness may internally signal the limits of applicability of the EFT framework. When we assume that theories are natural, we assume that EFT methods apply and that the effects of high-energy physics on low-energy Lagrangians are relegated to fixing the values of coupling constants.[8]

Wallace (2019) argues that far from being a technical requirement relevant only to high-energy physics, naturalness underwrites how we understand inter-theoretic relationships like emergence and reduction throughout physics. In Wallace's account, naturalness plays an essential role in deriving emergent dynamics for macroscopic systems from more fundamental theories.[9] We argue that the failure of naturalness in the CCP may signal the limits of applicability of the EFT framework. The EFT approach is overstated if taken to be a precondition for the possibility of physical theorizing, as one reading of Wallace (2019) suggests. Although we acknowledge the wonderful utility of decoupling, there is no necessity that nature cooperates with our fondness for EFTs. By rejecting the global applicability of the EFT framework, we endorse pursuing "unnatural" solutions.

The article proceeds as follows. Section 2 more carefully states the CCP within the EFT framework. We argue that there is no direct path to the CCP in terms of a conflict of differing measurements of Λ from different observations. Within the Standard Model, there is no evidential support for any particular value of vacuum energy density. Thus, the problem arises in the context of treating GR as an EFT and using the renormalization group to understand the scaling behavior of Λ. Yet unlike the scaling of other terms in the effective action, a shift in the value of Λ threatens to undermine assumptions about spacetime implicit in this way of treating the problem. Section 3 considers this question from a different perspective. We make explicit the spacetime structure that standard EFT techniques depend on, then examine the ways in which those spacetime assumptions can be relaxed for applications of GR as an EFT. The relaxed assumptions allow for EFT methods to be applied in special cases where the spacetime is nearly static or asymptotically flat. But the CCP

---

[7] Schneider (2020, 2022) has argued that the problem takes different forms depending on one's interpretative stance toward QFT and GR, suggesting different strategies to (dis)solve the problem in quantum gravity. Here, we aim to give a clear formulation of the problem within the EFT framework. We take a further step of arguing that the best solution to the problem is to dissolve it by rejecting (i) and (ii).

[8] See, in particular, Williams (2015, 2021), Wallace (2021, 2019), Rivat (2019); Rosalerand Harlander (2019); Ruetsche (2020) for philosophical discussions most closely related to our concerns.

[9] Wallace considers "naturalness" in broader terms than we will here, applying the notion to probability distributions in statistical mechanics and the emergence of time-asymmetric macroscopic dynamics, as well as to QFTs.

arises when considering large-scale features of the universe, and EFT methods break down in this regime. Thus, it should not be surprising that EFT methods fail for understanding Λ. The failure of decoupling serves as an internal signal that the approach fails, and the limitations of EFTs support this conclusion from an external perspective. In section 4, we discuss some approaches to quantum cosmology that fall outside the EFT framework. The purpose of this section is to illustrate that the EFT framework, decoupling, and naturalness are not necessary preconditions for constructing models in physics. Finally, section 5 returns to the question of naturalness and its necessity for doing physics.

2. The cosmological constant problem

We characterized the CCP as arising from treating GR as a low-energy EFT. But is there a more direct way of posing the CCP? For example, if we have direct evidence that vacuum energy $\langle\rho\rangle$ exists, and it should contribute to the Einstein field equations as an effective Λ term, doesn't this immediately lead to a conflict—that different ways of inferring the same quantity lead to wildly different results? We deal with this question in section 2.1, concluding that there is no independent evidence for $\langle\rho\rangle$ from the point of view of QFT. We therefore have a problem with the EFT formalism when extended globally, as we indicate in section 2.2.

2.1 No conflicting measurements of Λ

Consider the effective Einstein–Hilbert action coupled to matter in the form of quantum fields (eq. [1]). The stress-energy tensor for matter fields will include a vacuum energy density playing an analogous role in the Einstein field equations to the cosmological constant. In semiclassical form, this looks like

$$\langle T_{ab}\rangle = \langle\rho\rangle g_{ab}, \tag{3}$$

where the expectation value is taken about the global vacuum state. Because both Λ and $\langle\rho\rangle$ contribute as constant multiples of the metric, we observe only the consequences of their combination,

$$\Lambda_{\text{obs}} = \Lambda + \kappa\langle\rho\rangle. \tag{4}$$

If we have direct evidence for the presence and value of $\langle\rho\rangle$ in $\mathcal{L}_m$, then it should contribute to Λ_{obs} along with the Λ term from the Einstein–Hilbert action. This apparently allows for a direct observational comparison: measure the total energy density in a region, including $\langle\rho\rangle$, and compare it to the curvature revealed through cosmological observations. However, any such "prediction" of $\langle\rho\rangle$ has to resolve ambiguities associated with composite operators (polynomials of field operators) in interacting QFTs. Here, we will focus, in particular, on ambiguities regarding the stress-energy tensor.

In perturbative QFT, the field operators appearing in a Lagrangian have no direct physical significance: we can write the Lagrangian in terms of new fields. When dealing with renormalizable QFTs, the only possible redefinitions are linear transformations, whereas EFTs allow for integer polynomials of ϕ and a finite number of derivatives. Such field redefinitions do not change the S-matrix elements. Physicists have taken advantage of this freedom to remove divergences by expressing

the Lagrangian in terms of renormalized fields.[10] Further natural constraints are imposed to clarify the physical meaning of some operators; for example, in the case of a conserved current J^μ associated with an internal symmetry, there is no ambiguity in defining the operator (Collins 1985, §6.6). The stress-energy tensor T_{ab} includes products of field operators. For any of the methods introduced to handle these products, we can ask whether they rule out a field redefinition that has the following impact on the stress-energy tensor: $T'_{ab} = c_0 T_{ab} + c_1 \eta_{ab} \mathbf{I}$ (where $\mathbf{I}$ is the identity operator). Redefinitions in the EFT approach are typically required to preserve S-matrix elements and n-point functions. It turns out that preservation of the S-matrix does not constrain the value of c_1 because the total energy cancels out in calculations of the S-matrix elements. Thus, it does not appear that QFT has the resources to predict an unambiguous value for vacuum energy density.

Nevertheless, articles on the cosmological constant abound with claims that QFT predicts a value of vacuum energy density. For the sake of argument, consider the current best estimates,[11] $\langle \rho \rangle \simeq -2 \times 10^8 \text{GeV}^4$, differing by over 50 orders of magnitude from the value of $\Lambda_{\text{obs}}/\kappa$ fixed by cosmological observations, $\simeq 10^{-47} \text{GeV}^4$. We need not appeal to cosmology: even solar system dynamic constrains $\Lambda_{\text{obs}}/\kappa$ to be ≈ 40 orders of magnitude smaller than $\langle \rho \rangle$.

The attempt to directly relate gravitational measurements of Λ to the vacuum energy density requires two assumptions. The first is that $\langle \rho \rangle$ gravitates. One class of the *modified dynamics* approaches to solving the CCP rejects this. By introducing a mechanism or modification that decouples $\langle \rho \rangle$ from gravity, one can treat $\Lambda_{\text{obs}} = \Lambda$ as a free parameter, determined by observation. For now, we will assume that vacuum energy, if real, obeys the equivalence principle like all other forms of energy. If not, then we still do not arrive at conflicting measurements of the same quantity; in that case, $\langle \rho \rangle$ does not contribute to Λ. The second assumption is that the vacuum expectation value of energy density is real, that is, not an artifact of the QFT formalism on Minkowski spacetime. Its value must be determined independently of considerations of gravity; otherwise, the input $\langle \rho \rangle$ is unknown. Do we have direct evidence for the reality (and magnitude) of $\langle \rho \rangle$ from the Standard Model? How seriously should we take predictions of its value, such as the one cited earlier?

The standard response is to claim that either the Lamb shift or the Casimir effect provides direct evidence of the presence of vacuum energy density. However, both effects, at best, provide evidence for the presence of local fluctuations in vacuum energy, not a global expectation value (Koberinski, 2021a). Typically, the Casimir effect, described as due to impenetrable plates limiting vacuum fluctuation modes, is taken as the strongest evidence in favor of $\langle \rho \rangle$. The plates constrain the production of virtual photons in the vacuum—only photons with wavelengths that are an integer

---

[10] For example, in the case of QED, the electron and photon fields are multiplied by renormalization factors Z_2, Z_3: $\psi_0 = Z_2^{1/2} \psi$ and $A_0^\mu = Z_3^{1/2} A^\mu$. Such field redefinitions lead to new Green's functions, but the S-matrix elements will be the same as long as $\langle p | \theta | 0 \rangle \; /= 0$ (i.e., the field can create a one-particle state with momentum p from the vacuum).

[11] This is the value that Martin (2012) obtains using dimensional regularization, along with renormalizing using modified minimal subtraction at first order. This value hides the sensitive dependence on higher-mass scales that are reintroduced when higher-order terms are considered. Koberinski (2021b) argues that this quantity does not meet the standards of a candidate prediction. We set this issue aside for now to focus on the lack of direct evidence for $\langle \rho \rangle$.

multiple of the plate spacing can be created between the plates. This creates a pressure differential because "more" virtual photons can interact with the outside of the plates than the space in between, leading to a small attractive force. However, alternative formulations characterize it as a residual van der Waals force between the atoms in the plates; Jaffe (2005) has explicitly performed an alternative calculation in which the effect is due to loop corrections in the relativistic forces between the material plates. This calculation generalizes more readily to other plate geometries, and unlike a pure vacuum pressure, it goes to zero when the QED coupling α is taken to zero. The original explanation in terms of differential vacuum pressure may be a successful shorthand for the more realistic explanation, but it seems to be little more than that.

For the Lamb shift, it is even clearer that this is nothing more than radiative corrections to a first-order QED calculation. The Lamb shift is a small difference in the $2s$ and $2p$ orbital energy levels of the hydrogen atom, which are equal if one uses the Dirac equation. From QED, we see the effect as a one-loop correction to the interaction between the proton and electron in a hydrogen atom. Loop corrections to interactions are not the same as vacuum energy, even if they are sometimes fancifully described as virtual particles from the vacuum interacting with the external particles. At best, these should be thought of as quantum fluctuations about the vacuum state. In terms of Feynman diagrams, vacuum energy is represented as a sum of bubble diagrams—diagrams with no external legs. These diagrams factor out of any n-point function and therefore play no role in predictions based on perturbation theory.

To summarize the arguments of this section, we claim that $\langle \rho \rangle$ plays no role in the empirical success of the Standard Model and that, furthermore, the Standard Model provides no prediction of its value. We cannot generate a direct conflict between different ways of measuring Λ. Instead, we must deal directly with the principles of EFT for cosmological spacetimes.

2.2 The cosmological constant in effective field theory

The fundamental quantities of a QFT are the correlation functions among a set of operators $\{O_i\}$ acting on the vacuum state, calculated based on the action $S = \int \mathrm{d}^4x \mathcal{L}(\phi)$ for a specific field theory (schematically):

$$\langle O_1, \ldots O_n \rangle = \int \mathcal{D}\phi \exp^{iS(\phi)} O_1(\phi) \ldots O_n(\phi). \tag{5}$$

The EFT approach deals directly with these quantities, explicitly indexing them to a particular energy scale. Because the action is now defined in terms of effective degrees of freedom at that energy scale, we think of it as an effective action for that domain.[12] This effective action can be constructed "top down" from an existing high-energy theory, such as by systematically integrating out the high-energy degrees of freedom, given a cutoff scale Γ. This can be described more abstractly as the action of the renormalization (semi-)group on the space of theories, that is, actions at specific energy scales $\{S(\Gamma)\}$. This group generates a trajectory relating actions at different

---

[12] For reasons of space, we cannot go into detail regarding the EFT framework. For pedagogical overviews, which we draw on here, see Manohar (2020); Donoghue (2012); Burgess (2004), and see Williams (2021); Ruetsche (2020) for introductions for philosophers.

scales, and in the best case, trajectories through the infinite-dimensional space of theories $\{S(\Gamma)\}$ flow to a finite-dimensional subspace.

EFTs constructed "top down" in this fashion, from a given high-energy theory, provably yield low-energy observables compatible with the results of the full theory. We can also develop an EFT "bottom up"—proposing a Lagrangian $\mathcal{L}_{\text{eff}}$ with appropriate symmetries and fields, and including all possible couplings consistent with those symmetries, even though it is not obtained from a known high-energy theory. A separation of scales is still needed in the bottom-up approach. Obviously, one cannot then prove directly that the EFT will approach the (unknown) high-energy theory. The absence of the high-energy theory means that in applying the EFT framework, we must make substantive assumptions about an unknown future theory. One of these assumptions is clearly locality, namely, that $\mathcal{L}(\phi_i)$ depends on the fields ϕ_i and their Taylor expansions at a point.

In the EFT framework, we can classify the behavior of the vacuum energy density under renormalization group flow. To see why decoupling fails for a vacuum energy density term, we must first explain the behavior of different terms in the Lagrangian. In a spacetime with four dimensions,[13] couplings with positive mass dimension indicate relevant parameters that increase in magnitude in the EFT as the cutoff is taken to higher energies. Renormalizable theories contain these and marginal parameters in the Lagrangian, the latter characterized by dimensionless couplings, which therefore do not contain powers of the cutoff. Irrelevant terms have coupling constants with dimension of negative powers of mass. Decoupling applies to the marginal and irrelevant parameters; relevant terms appear to couple sensitively to the high-energy cutoff. A sensitive dependence on the cutoff signals that relevant terms are sensitive to the scales at which new physics comes in.[14]

The vacuum energy density $\langle \rho \rangle$ and Λ terms exhibit the most problematic scaling behavior: they are relevant parameters, and they scale with the fourth power of the cutoff. The Standard Model is well confirmed up to the energy scales probed so far at the Large Hadron Collider (LHC), so the cutoff for an effective version of the Standard Model must be at least 1 TeV. One can arrange a delicate cancellation between the scaling from vacuum energy density plus quantum corrections to GR and the bare Λ term: $\Lambda_{\text{obs}} = \mathcal{O}(\Gamma^4) - \mathcal{O}(\Gamma^4) \approx 0$, but this seems ad hoc. Further, it is unstable against radiative corrections to the vacuum energy density obtained when the higher-order terms in a perturbative expansion are included.[15] Because we do not

---

[13] The number of spacetime dimensions applies here to the *effective* theory. If GR can be treated as a low-energy effective limit of string theory, for example, it must be possible to approximate the relevant domains of string theory after the compactification of extra dimensions.

[14] Although this more literal reading of the significance of the cutoff is common among physicists, it is not uncontroversial—Koberinski (2016, 2021a, 2021b) critiques an overly literal interpretation of the cutoff scale in EFTs; Rosaler and Harlander (2019) argue that all theories in the equivalence class of theories related by renormalization group transformations are actually the same theory. Until new fields are introduced, "the same" EFT with a different cutoff scale leads to the same predictions. Here, we set these issues aside and illustrate the CCP using the more standard view of EFTs.

[15] As noted in other approaches (cf. Martin (2012)), dimensional regularization provides a better grounds for formulating the CCP in the EFT framework. In that case, $\langle \rho \rangle$ depends on the fourth power of the masses of all fields in the Standard Model. We state here the "standard" EFT account depending on a cutoff regularization, although the issue persists under dimensional regularization.

observe $\langle \rho \rangle$ directly, this is not an empirical problem. It instead indicates a breakdown of decoupling within the EFT framework. The behavior of $\langle \rho \rangle$ under renormalization group flow suggests that vacuum energy density is sensitive to high-energy physics. If the local, relevant $\Lambda + \langle \rho \rangle$ term from equation (1) is extrapolated to provide a contribution to the observed cosmological constant, this would indicate a highly sensitive coupling between high-energy physics and the deep infrared (IR) in cosmology.

There is a further challenge regarding how to understand this scaling behavior: Is there a self-consistent choice of background metric (and other structures) we can use to describe the renormalization group flow? Suppose that we start with an action defined at a specific energy scale E_h, low enough so that quantum gravity effects can be neglected and the metric g_{ab}^1 is a solution of classical GR. Implicitly relying on this metric, we can integrate out the high-energy modes to obtain an effective action at a lower-energy scale E_l. Yet the appropriate metric *cannot* still be g_{ab}^1 at this lower scale because the scaling properties described previously lead to a nonzero Λ contribution. Even for relatively small changes of scale, this term will dominate, such that the action at the lower scale is defined with respect to a different classical metric, g_{ab}^2. It is common to claim that generic curved spacetimes "look enough like Minkowski locally," such that the tools developed in flat space can be used. But incorporating the scaling of vacuum energy leads from an initial spacetime g_{ab}^1 to one with strikingly different global properties—for example, from Minkowski spacetime to de Sitter spacetime. Minkowski spacetime is qualitatively different from de Sitter spacetime, no matter how small the value of Λ. Furthermore, the $\Lambda \to 0$ limit is not continuous, as illustrated by the contrast in conformal structure. This suggests that the renormalization group trajectory for Λ should be defined over a space of metrics, not just over the values of parameters appearing in the Lagrangian.

In sum, the scaling behavior of Λ within the EFT approach signals a dependence on high-energy physics, and we have argued that it also cannot be consistently described with respect to a single fixed background spacetime. This raises the broader question of what we need to assume regarding spacetime to apply EFT techniques, which we turn to next.

3. Spacetime for effective field theories[16]

Effective field theories, as generalized from renormalizable QFTs, implicitly rely on certain background spacetime structures. Both top-down and bottom-up construction procedures partition the degrees of freedom for a system into those relevant to the EFT and those outside of its domain. Typically, the EFT describes low-energy, fluctuating modes against a backdrop of high-energy modes that remain in an adiabatic ground state. Such a description relies on separating the degrees of freedom based on their energy, which requires a well-defined notion of energy, as well as a sufficiently stable cutoff point to sort high- from low-energy degrees of freedom. This means that the spacetime on which the EFT is defined must have something approximating a timelike killing vector field. This is a demanding requirement, not satisfied by, for

---

[16] We are grateful to an anonymous reviewer for detailed comments that prompted revisions to this section.

example, the FLRW models used in relativistic cosmology. This does not threaten the insights gained from treating GR as an EFT, applied to problems that assume either a Minkowski background or some other background with sufficient structure (at least approximately) to identify the relevant degrees of freedom. Yet it does raise the question of how much insight we can gain from EFT methods regarding the cosmological constant.

Applications of EFT methods proceed, schematically, by identifying low-energy degrees of freedom and symmetries, then writing the most general effective action for these degrees of freedom compatible with these symmetries. The earlier form of the effective action (eq. [2]) follows by treating the low-energy degrees of freedom as gravitons (spin-2 fields), along with matter degrees of freedom, and requiring diffeomorphism invariance and local Lorentz invariance for terms in the expansion. There are several other ways of applying EFT techniques to gravitational physics, such as the "nonrelativistic GR" approach (Goldberger and Rothstein 2006) developed to study the in-spiral phase of merging compact objects and the radiation they emit.[17] This approach takes advantage of the separation of scales between the extended compact objects and gravitational perturbations, integrating out the degrees of freedom associated with the objects and treating them as point particles, and starts from a different effective action. EFT techniques have also been applied to the study of structure formation in cosmological models, based on an action that describes a coupled scalar field–metric system satisfying the FLRW symmetries.[18]

Here, we will focus on an EFT constructed for gravity based on the effective action given in equation (2). EFT calculations based on this action have led to several seminal results, such as Donoghue's expression for the leading-order quantum corrections to the Newtonian potential between nonrelativistic particles.[19] The higher-order terms in the Lagrangian scale with the inverse powers of the Planck mass, so the quantum corrections are extremely small. Because there is a much larger separation of scales here than in other areas of physics, the EFT for GR is sometimes described as the best EFT. Yet the cosmological constant is not dynamically relevant in this calculation, which proceeds in Minkowski spacetime. Donoghue (2012), for example, explicitly treats Λ as one of the EFT parameters to be fixed by observations, and he simply sets it to zero in calculating the quantum corrections while noting that it is unimportant in this domain. As we will see, this is only permissible when we have an external reason to think that Λ is not physically relevant.

Extending beyond Minkowski spacetime, it is still necessary to identify the degrees of freedom to be included in the action and draw the contrast between high- and low-energy modes. As we noted earlier, this is possible in static spacetimes with a timelike Killing vector field. In static spacetimes, we have a well-defined separation of energy scales—and therefore a well-defined notion of energy conservation—and can identify a stable ground state. Furthermore, we can construct a conserved energy that is

---

[17] This is called a "nonrelativistic" approach because the relative velocities of the two objects are small during this phase; EFT techniques have been used to streamline and extend results that had been calculated earlier using a post-Newtonian perturbative expansion.

[18] In addition to proceeding from a different action, in this case the EFT methods have been extended to cover "open" systems due to the migration of modes across length scales during inflation (see, e.g., Burgess 2017).

[19] See Burgess (2004); Donoghue (2012) for overviews and references to the original literature.

bounded from below. This naturally gives rise to a well-defined vacuum state as the lowest-energy eigenstate of the Hamiltonian operator. In general, a frequency splitting for matter fields can be carried out as well. Given all of this, we can identify perturbations around this vacuum and create a Fock space of fluctuations, and also distinguish between low- and high-energy states, in order to apply EFT methods.

Physicists have successfully applied EFT techniques to spacetimes that have approximately static regions and those that are symmetric "at infinity" (i.e., quasi-static and asymptotically flat spacetimes, respectively; see Burgess [2004] for an overview). For the former, as long as a local, approximate notion of energy remains well defined for the timescales relevant to the problem, one can construct an *approximately* conserved Hamiltonian and create an *approximate* division into high- and low-energy modes. But these relaxed conditions still depend on an *approximately* well-defined separation of energy scales. For backgrounds on which these approximations fail for the distance and timescales of interest, the EFT construction procedure cannot get off the ground.

In asymptotically flat spacetimes, one can exploit the Minkowskian structure at infinity to define conserved energies and ground states. Provided that one is interested in effects observable far away from the central region with complex gravitational dynamics, it is reasonable to expect that EFTs provide a good basis for calculation. This is the assumption behind EFT calculations of Hawking radiation measured far from the event horizon of a black hole.

One generalization most relevant to the domain of cosmology is that to slowly varying time-dependent background spacetimes. In general, one cannot construct an EFT without energy conservation because EFTs organize and separate states according to energy. However, if the time evolution is adiabatic—that is, the metric and other time-dependent fields vary sufficiently slowly compared to the ultraviolet (UV) scales of interest—one can construct an approximately conserved Hamiltonian, an approximate ground state, and an approximate (time-dependent) low-/high-energy split (cf. Burgess (2017)). Adiabatic evolution is then indexed to particular domains of a spacetime solution. Where adiabaticity fails, one can encounter crossing of energy scales, from the EFT $p < \Lambda(t)$ to the high-energy regime $p > \Lambda(t)$, and vice versa.

So far we have focused on the spacetime structure needed to identify the degrees of freedom of interest, to take the first step in constructing an EFT. But the full force of the EFT framework provides more than just this first step. It is not enough to cut off high-energy degrees of freedom; we must also ensure that the resulting theory has the appropriate interscale insensitivity using renormalization group methods. Without this assurance, we have effectively fine-tuned a solution that is neither self-consistent nor robust to perturbations at energy scales that are supposed to have been screened off. We can think of this as a two-stage process for setting up EFTs. First, can we write down an EFT at a given scale, setting its couplings to those determined empirically, and use this to calculate leading-order quantum effects on gravity? We argue that the answer to this question is yes: the successful applications of EFTs in gravity described earlier take this form. Second, can we then extend this EFT to different (higher- or lower-) energy scales, using the scaling properties typical of flat-space EFTs? The answer to this question is no: the failure of naturalness for

Lambda ruins the possibility of a self-consistent background metric that grounds the notion of energy and that of scale separation for the EFTs at different energies.

The scaling behavior discussed at the end of the previous section raises a different challenge. The EFT describes low-energy degrees of freedom propagating with respect to a fixed background, such as a vacuum solution or thermal state. Can we assume that the background used in the EFT is consistent with the solutions to field equations of the full theory? This can be proven to hold (in cases where the full theory is known) if the background fields evolve adiabatically (see Burgess 2017). But the background fields do not remain static if quantum corrections to the action have the form of an effective cosmological constant term. As discussed in section 2.2, a spacetime background upon which an EFT is constructed will change drastically under an otherwise standard scaling transformation of the EFT, undermining the self-consistency of the full EFT treatment. The backreaction of such contributions on the metric may be negligible in relatively small spacetime regions, but they have a cumulative effect at large distances and over long times. Hence, it is strikingly implausible to assume that the EFT background matches a solution to the full-field equations at large scales. Yet these are precisely the scales at which the dynamical effects of a cosmological constant term would become apparent.

These backreaction effects can take a different form in the deep IR as well. Assume that we can model a "patch" of a given spacetime using either Minkowski spacetime or some other fixed background spacetime. In that case, we can use Riemannian normal coordinates in a local patch, and we can make it clear in what sense the spacetime "looks locally Minkowskian." Thus, within that patch we have a well-defined background on which to construct an EFT. We can similarly construct local patches over other regions of the spacetime. However, to be able to stitch these together, we would need to impose a strong constraint on the metric (or on the curvature) that is not likely to hold in general. Donoghue (2009) takes this situation as a novel illustration of "how EFTs fail," in that they cannot adequately describe the "buildup" of effects as we patch such local descriptions together.

Both points challenge the assumption that we have adequate control over the background solution to establish the self-consistency of the EFT. The consistency challenge arises if we try to maintain both: (i) we will carry out EFT calculations describing spin-2 and matter fields propagating over Minkowski spacetime (or in a static, curved spacetime), and (ii) the quantum fields contribute to a nonzero cosmological constant as a result of radiative corrections. If (ii) is accepted, then the assumption that the spacetime background is Minkowski (or static) is, at best, an approximation for limited regions. The appropriate background at larger scales—if one can even be defined—should instead be de Sitter spacetime because of the large Λ contribution. When looked at internally, we arrive at the reductio of the CCP. Externally, we have shown that one should not expect EFT methods to get off the ground in spacetime settings without approximate or asymptotic temporal symmetries relative to the physics of interest. Spacetimes where Λ is dynamically relevant, including FLRW cosmological models, fall outside the domain of current EFT treatment. New conceptual resources are needed when it is relevant. We outline some potentially promising avenues in the next section.

4. Unnatural solutions

As Kuhn (1962) recognized, criticisms of an appealing approach rarely lead scientists to abandon it unless there is an available alternative. We have argued earlier that the EFT program is ill-suited for dealing with global features of cosmology and, in particular, that the CCP is a signpost that something has gone wrong. Although there is not yet a clear alternative, there are several lines of work that aim to reformulate the foundational principles of flat-space QFT. These avenues of research show that the separation of energy scales, far from being a precondition for the possibility of science, is not an essential feature of current speculative physics. To be clear, we do not expect the EFT approach to be entirely replaced, given successes such as the EFT methods applied to GR mentioned earlier. Rather, the CCP forces us to acknowledge the limitations of the EFT approach, alongside the need for new ideas regarding the global properties of quantum fields coupled to gravity. In this section, we briefly outline three research programs that reject some of the basic EFT concepts: quantum field theory on curved spacetimes, the UV-IR correspondence, and the breakdown of locality from string theory. Some of these approaches reject the EFT framework in the context of matter fields on classically curved spacetime backgrounds, whereas others reject it directly for gravitational degrees of freedom. In either case, these research programs highlight the ways in which the decoupling of energy scales fails in the cosmological solutions relevant to the problem at hand.

QFT on curved spacetimes

This approach replaces foundational concepts in QFT with generalized versions appropriate for generic curved spacetime backgrounds. It takes to heart lessons from GR in aiming to construct quantum field theories in a way that depends only on local spacetime properties. The spacetime background is still treated classically, with the generalizations focused on the equations governing matter fields. Conceptual reen-gineering focuses on the spectrum condition, Poincaré covariance, and the existence of a unique vacuum state because all of these depend on or follow from symmetries of Minkowski spacetime. The ambitions of this approach do not extend to including backreaction of the quantum fields on spacetime; this is not a quantum theory of gravity, and it is contentious how much it contributes to formulating one.[20] Essentially, this approach deals only with understanding matter degrees of freedom and ignores the treatment of gravity as an EFT. Nonetheless, it highlights one way to make fundamental changes to our understanding of QFT, as well as the resulting effects on the EFT framework and the CCP.

For the sake of definiteness, we focus here on the axiomatic approach pursued by Hollands and Wald (for reviews, see 2010, 2015), based on constructions of simple scalar ϕ^4 models on globally hyperbolic, but otherwise generically curved, spacetimes. They work in a position-space representation and use operator product expansions as the basic local building blocks for a QFT, rather than the Fock-space momentum

---

[20] In favor of this approach, one might argue that a classical background spacetime can be treated as a coherent state of gravitons, then treat the quantum fields as light enough to have a negligible effect on spacetime curvature (cf. Dvali, Gómez, and Zell [2017] and references therein). Whether this assumption is a sensible one depends in part on its consistency. For the purposes of our argument, however, it is only important to see how QFTs on curved spacetime revise key concepts that lead to the CCP.

representation on which Minkowski QFTs are built. Although a Fock-space representation for free fields is not necessarily required for a generic EFT, one often does assume that many of the symmetries of Minkowski spacetime hold locally. Moreover, if transition amplitudes are meant to be transitions between privileged, well-defined particle states, then the Fock-space construction is required. Hollands and Wald abandon this framework: Poincaré covariance is generalized to a local general covariance of the fields, and the positive-frequency condition for fields is characterized locally in terms of the singularity structure of the n-point functions of fields. This is the microlocal spectrum condition: it encodes the same information as the positive-frequency condition, but it does so in a local way that does not depend on the global structure of spacetime.

The key conceptual change is the lack of a vacuum state as the basis for constructing a QFT. On curved spacetime backgrounds, there is no privileged global vacuum state, and therefore nothing that has the correct symmetries and invariance properties to play the role of a cosmological constant term once gravitational degrees of freedom are included. This lack of a preferred vacuum marks a sharp contrast with conventional QFT, which often aims to calculate correlation functions for quantum fields in their vacuum states. Furthermore, renormalization techniques in flat spacetime implicitly utilize the preferred vacuum state in order to handle products of field operators. These renormalization procedures are often presented as subtracting divergences mode by mode. By contrast, Hollands and Wald's local and covariant formulation of QFT in curved spacetimes has to do without a globally defined preferred state or a division into positive-/negative-frequency modes. The treatment of renormalization they develop is by necessity holistic (see §3.1 in Hollands and Wald 2015): products of field operators have to be renormalized with respect to a locally defined quantity (the Hadamard distribution), and this cannot be interpreted as a mode-by-mode subtraction. As a result, it is difficult to see how to implement the division between low- and high-energy modes that is crucial to EFT methods.

If this approach is used as a starting point for quantizing gravity, it is unclear how one would construct EFTs. The approximation of small perturbations about a static or asymptotically flat background fails for generic globally hyperbolic spacetimes, as we have argued in the previous section. This approach to QFTs on curved spacetimes illustrates one way of rethinking the foundations of QFT and the basic elements of renormalization when merging gravitational and matter degrees of freedom. By considering how curved spacetime backgrounds change the construction of QFTs, one is less tempted to inappropriately generalize the successes of the EFT framework to generic globally hyperbolic spacetimes.

Breakdown of naturalness from quantum gravity

Although there is not yet a satisfactory theory of quantum gravity, we can look to the most developed speculative theories to see the ways in which the EFT approach might break down. By focusing on insights from candidate theories, such as string theory, we can gain insight into the ways that low-energy physics might be affected by a future complete theory of quantum gravity. There are multiple ways that EFT methods could break down in any theory of quantum gravity. First, the successor theory might introduce new physics at some intermediate scale (below the Planck

scale) that naive EFT approaches would miss if they are taken to be applicable right up to the Planck scale. Second, the emergence of spacetime would place limitations on the applicability of EFTs. The assumption of a static or asymptotically flat spacetime background must break down when the very concepts of space and time also break down. In regimes where spacetime concepts fail to apply, the ideas of background spacetime, locality, and separation of energy scales also fail to apply.

These two generic "breakdowns" are of less interest because neither would require a fundamental reconfiguration of the EFT approach. In the former case, the general EFT methodology would still apply, and one would simply have to lower the upper limits of applicability of EFTs to the new mass scale. In the latter case, as long as all interactions in the successor theory are local—at scales where the concept of locality remains relevant—EFTs would be insensitive to the breakdown of spacetime at scales far below the breakdown.

One more interesting problem related to the latter involves theories of quantum gravity whose fundamental degrees of freedom are non-spatiotemporal. At a more fundamental level, one might worry that a global separation of energy scales makes no sense for the fundamental degrees of freedom. The conceptual understanding of integrating out high-energy degrees of freedom would then break down in light of the new theory because the concept of "high energy" may not be well defined. There is little reason to suspect that fundamentally non-spatiotemporal degrees of freedom can fit the EFT framework and that their effects will be limited to renormalizations of coupling constants or additional local interactions.

Other interesting problems arise when we consider specific theories of quantum gravity. String theory raises two potential problems with the EFT approach. First, string theoretic T-duality links the UV and IR energy scales. We save this until the next section because a UV-IR correspondence can arise in other contexts (e.g., double-field theory). The second problem is the breakdown of locality at the string scale. Because strings are extended objects, at length scales comparable to that of the strings, the idealization of treating string interactions as local point interactions breaks down. This may not be a problem in static or asymptotically flat spacetimes because the nonlocalities at the string scale are unlikely to have impacts at larger distances (ignoring the deep IR scales and T-duality). If string-scale interactions do not grow or cascade over time, then EFTs at much lower-energy scales can deal effectively with nonlocal string interactions the same way that any QFT does: by approximating the string length to be zero and treating the interactions as local. At energy scales much lower than the string scale, this approximation will hold, and the EFT approach should proceed without problems.

How do these considerations bear on extending EFT techniques beyond static or asymptotically flat spacetimes? In cosmological spacetimes, we need to ensure that there are no cascading effects across energy scales that would lead to a stretching of nonlinear effects originating at the string scale. One concrete example where this has been conjectured is in inflation. In a rapidly expanding universe, Planck- or string-scale fluctuations would stretch rapidly; if inflation went on for a long enough time, these fluctuations could cross the Hubble radius and classicalize. Bedroya et al. (2020) proposed a trans-Planckian censorship conjecture, ruling out by fiat the possibility of Planck-scale modes crossing the Hubble radius. The censorship conjecture limits the length of inflation to be short enough that Planck-scale fluctuations cannot grow to a

size comparable with the Hubble radius. Further, this only rules out Planck-scale physics; if nonlocalities arise at $l_{\text{string}} \ll l_{\text{Planck}}$, then the constraint on the length of the inflationary epoch is even more restrictive. The censorship conjecture is a patchy fix, and Bedroya et al., acknowledge that something beyond an EFT approach may be needed to properly address the problem. New ideas that can reproduce the power spectrum of anisotropies in the cosmic microwave background radiation may be needed for the early universe because inflation stretches EFT methods beyond their proper domain of applicability. Some alternatives to inflation inspired by string theory are the emergent universe from a string gas and an ekpyrotic bounce universe (Brandenberger 2014). Loop quantum cosmology also posits a bouncing universe. All of these are capable of producing a nearly scale-invariant power spectrum for anisotropies, and so they may be alternatives to inflation. Additionally, two of these approaches—loop quantum cosmology and string gas cosmology—also abandon the standard EFT framework. Instead of a bottom-up EFT construction, they start with the high-energy theory and work down to the appropriate limiting domains applicable to early-universe physics. Although inflation is by far the most well-developed approach to understanding structure formation in the early universe, these competitors show how one could move forward outside of the context of EFTs.

UV-IR correspondence

One final approach to speculative physics that rejects decoupling and the EFT framework is the idea of a symmetry or correspondence between high-energy (UV) and low-energy (IR) degrees of freedom. It is relatively obvious how a UV-IR correspondence would fall outside of the EFT approach: at very high and very low energies, physical effects are sensitively coupled, so EFTs cannot be used to integrate out the effects of high-energy physics, at least without knowing the exact form of the high-energy theory. A UV-IR correspondence could apply directly to gravitational degrees of freedom or to matter degrees of freedom. Typically, a UV-IR correspondence is discussed in relation to the T-duality symmetry in string theory, where there is a transformation between degrees of freedom at the distance scale R and its inverse $1/R$, where distance is in string units. In the case of the cosmological constant, one might explain its presence and particular observed value as a remnant from some high-UV effects because cosmological distance scales are in the deep IR. From the point of view of string theory, T-duality has implications for both matter and gravitational degrees of freedom, and it therefore has the potential to link the two with the cosmological constant.

In one sense, the EFT framework is capable of accommodating T-duality because it is strictly another symmetry from the high-energy theory. One must include in their EFT description a spacetime symmetry mapping $x \to 1/x$ in the appropriate units. This is the project pursued by double-field theory. Essentially, one constructs an EFT with dual copies of the fields of interest, such that the duality is a symmetry of this expanded-field theory. Although this approach still falls within the EFT framework, naturalness is clearly violated, and it is unclear how the cosmological constant can arise in the context of double-field theory (Aldazabal, et al. 2013). Under the toroidal compactification approach in double-field theory, setting the flux through the torus equal to zero leads to vanishing scalar potentials and therefore no

cosmological constant term. However, more complicated compactifications may allow for a scalar potential to play the role of a cosmological constant term. Even though this approach leads to a type of EFT, it is of a very different character from the standard EFTs defined on a single configuration space. In particular, the decoupling assumption is violated because high energies in one set of degrees of freedom correspond to low energies in the other. Other approaches to handling a UV-IR correspondence may make further, more drastic modifications to the standard EFT approach.

5. Conclusions

The cosmological constant problem should be regarded as a reductio ad absurdum. As such, it shares the frustrating feature of any reductio argument, in that it shows that the set of starting assumptions leads to an absurd result without indicating where the error lies. Out of the assumptions generating the problem, we have argued that one should reject the application of EFT methods to the far IR of cosmological spacetimes. Standard EFT methods depend on having specific types of background spacetime structure. Although some generalization beyond Minkowski and Euclidean spacetimes is possible, EFTs cannot yet be constructed on generic spacetimes. As shown in section 3, EFTs require global properties to make a meaningful split between relevant and irrelevant scales, near static backgrounds to ensure this boundary doesn't change, and minimal backreaction effects on the spacetime background. None of these conditions holds for almost-FLRW spacetimes over cosmic scales of distance and time. The successes of treating GR as an EFT are limited to special cases where the background spacetime has sufficient structure, where backreaction effects are negligible, and where the cosmological constant can be ignored.

In general, induction from the success of limited examples to a greater scope of applicability is a good strategy for scientific inquiry. The great success of EFT methods might lead one to suspect that the issues with generalizing to curved spacetimes are merely transient and that the EFT methodology should not be abandoned. This is the approach that many working on solutions to the CCP have taken, implicitly or explicitly, and corresponds to the *modify-dynamics* approach outlined in the introduction. This is not an unreasonable approach, and it takes the reductio seriously by making local modifications to particular dynamical theories, rather than our proposal to reject the applicability of the EFT framework. Some solution strategies involve adding new symmetries (e.g., supersymmetry, double-field theory), whereas others involve adding new fields relevant at scales just above the Standard Model or modifying the coupling between gravity and vacuum energy. In most of these, one rejects decoupling and therefore concedes that some low-energy phenomena are tightly linked to high-energy physics. If one is confident in the EFT methodology, local modifications are still necessary to solve the problem.

However, the arguments in section 3 show the intrinsic limitations of the EFT framework and should temper any expectations that generalizations to cosmology will be successful. It could turn out that EFT methods will generalize, but such a generalization will require significant conceptual modifications to the basic assumptions, such as a global separation of energy scales. The key assumption that fails for the standard EFT approach here is the assumption of naturalness, in the form of

autonomy of scales. We have seen that the cosmological constant is highly sensitive to the choice of regularization procedure and the value of the regulator, indicating that it is sensitive to the details of high-energy physics. The apparent failure of naturalness in the CCP is well known, and it is often taken as part of a more general issue in physics. Besides the CCP, naturalness issues arise in the hierarchy problem in particle physics. Some argue that any failure of naturalness will have wide-ranging consequences for the metaphysics and epistemology of science generally (see Wallace [2019] and references therein). We have argued that such sweeping conclusions are unwarranted. Section 4 highlights some approaches that accept a local failure of naturalness within cosmology, rather than attempting to restore it or make radical changes to our understanding of science. Although a post-naturalness high-energy physics creates significant new theoretical challenges, we think this is an avenue worthy of serious pursuit.[21]

We end by noting that the separation of scales fails elsewhere in physics, and therefore the naturalness problems arising in gravitational EFTs are not sui generis. Nonlinear dynamical systems can exhibit inverse energy cascades, with energy transfer from short scales to longer scales—such as in turbulent motion. Such features would spoil the separation of scales essential to applying EFT methods. Mesoscale modeling also requires significant input from physics at distinct scales. Batterman (2013) has argued extensively in favor of mixing micro- and macro-scale modeling techniques in condensed matter physics and materials science, for example. Without accounting for physical effects at distinct scales, one misses important features of bulk materials. These phenomena are certainly more challenging to model and predict than those in which scales evolve autonomously. We can take some comfort in recognizing that this kind of challenge is not unique to quantum gravity and that physicists have developed techniques to handle the failure of scale separation effectively in other contexts.

References

Aldazabal, Gerardo, Diego Marques, and Carmen Nunez. 2013. "Double Field Theory: A Pedagogical Review." *Classical and Quantum Gravity* 30 (16):163001.

Batterman, Robert. 2013. "The Tyranny of Scales." In *The Oxford Handbook of Philosophy of Physics*, edited by Robert Batterman, 255–86. Oxford: Oxford University Press.

Bedroya, Alek, Robert Brandenberger, Marilena Loverde, and Cumrun Vafa. 2020. "Trans-Planckian Censorship and Inflationary Cosmology." *Physical Review D* 101 (10):103502.

Bianchi, Eugenio, and Carlo Rovelli. 2010. "Why All These Prejudices against a Constant?" *arXiv preprint.* DOI: 10.48550/arXiv.1002.3966.

Brandenberger, Robert. 2014. "Do We Have a Theory of Early Universe Cosmology?" *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 46 (part A):109–21.

---

[21] We are grateful to an anonymous reviewer for comments that prompted revisions to this paragraph.

Burgess, Cliff P. 2004. "Quantum Gravity in Everyday Life: General Relativity as an Effective Field Theory." *Living Reviews in Relativity* 7 (1):5.

Burgess, Cliff P. 2017. "Intro to Effective Field Theories and Inflation." arXiv preprint. https://arxiv.org/pdf/1711.10592.pdf.

Collins, John C. 1985. *Renormalization: An Introduction to Renormalization, the Renormalization Group and the Operator-Product Expansion.* Cambridge: Cambridge University Press.

Donoghue, John F. 2009. "When Effective Field Theories Fail." arXiv preprint. https://arxiv.org/ftp/arxiv/papers/0909/0909.0021.pdf.

Donoghue, John F. 2012. "The Effective Field Theory Treatment of Quantum Gravity." In *AIP Conference Proceedings*, vol. 1483, edited by Waldyr Alves Rodrigues Jr., Richard Kerner, Gentil O. Pires, and Carlos Pinheiro, 73–94. Woodbury, NY: AIP Publishing.

Duncan, Anthony. 2012. *The Conceptual Framework of Quantum Field Theory.* Oxford: Oxford University Press.

Dvali, Gia, César Gómez, and Sebastian Zell. 2017. "Quantum Break-Time of de Sitter." *Journal of Cosmology and Astroparticle Physics* 2017 (6):028.

Friederich, Simon. 2018. "The Asymptotic Safety Scenario for Quantum Gravity—an Appraisal." *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 63:65–73.

Goldberger, Walter D., and Ira Z. Rothstein. 2006. "Effective Field Theory of Gravity for Extended Objects." *Physical Review D* 73 (10):104029.

Goroff, Marc H., and Augusto Sagnotti. 1986. "The Ultraviolet Behavior of Einstein Gravity." *Nuclear Physics B* 266 (3–4):709–36.

Hollands, Stefan, and Robert M. Wald. 2010. "Axiomatic Quantum Field Theory in Curved Spacetime." *Communications in Mathematical Physics* 293 (1):85.

Hollands, Stefan, and Robert M. Wald. 2015. "Quantum Fields in Curved Spacetime." *Physics Reports* 574:1–35.

Jaffe, Robert L. 2005. "Casimir Effect and the Quantum Vacuum." *Physical Review D* 72 (2):021301.

Koberinski, Adam. 2016. "Reconciling Axiomatic Quantum Field Theory with Cutoff-Dependent Particle Physics." http://philsci-archive.pitt.edu/12496/.

Koberinski, Adam. 2021a. "Problems with the Cosmological Constant Problem." In *Philosophy Beyond Spacetime*, edited by Christian Wüthrich, Baptiste Le Bihan, and Nick Hugget, 260–79. Oxford: Oxford University Press.

Koberinski, Adam. 2021b. "Regularizing (away) Vacuum Energy." *Foundations of Physics* 51 (1):1–22.

Kuhn, Thomas. 1962. *The Structure of Scientific Revolutions.* Chicago: University of Chicago Press.

Manohar, Aneesh V. 2020. "Introduction to Effective Field Theories." In *Effective Field Theories in Particle Physics and Cosmology: Lecture Notes of the Les Houches Summer School: Volume 108, July 2017*, edited by Sacha Davidson, Paolo Gambino, Mikko Laine, Matthias Neubert, and Christophe Salomon, 47–136. Oxford: Oxford University Press.

Martin, Jerome. 2012. "Everything You Always Wanted to Know about the Cosmological Constant Problem (but Were Afraid to Ask)." *Comptes Rendus Physique* 13 (6–7):566–665.

Rivat, Sébastien. 2019. "Renormalization Scrutinized." *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 68:23–39.

Rosaler, Joshua, and Robert Harlander. 2019. "Naturalness, Wilsonian Renormalization, and 'Fundamental Parameters' in Quantum Field Theory." *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 66:118–134.

Ruetsche, Laura. 2020. "Perturbing Realism." In *Scientific Realism and the Quantum*, edited by Steven French and Juha Saatsi, 293–314. Oxford: Oxford University Press.

Rugh, Svend E., and Henrik Zinkernagel. 2002. "The Quantum Vacuum and the Cosmological Constant Problem." *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 33 (4):663–705.

Schneider, Mike D. 2020. "What's the Problem with the Cosmological Constant?" *Philosophy of Science* 87 (1):1–20.

Schneider, Mike D. 2022. "Betting on Future Physics." *British Journal for the Philosophy of Science* 73 (1):161–83.

Wallace, David. 2019. "Naturalness and Emergence." *The Monist* 102 (4):499–524.

Wallace, David. 2021. "The Quantum Theory of Fields." In *The Routledge Companion to the Philosophy of Physics*, edited by Eleanor Knox and Alastair Wilson, 275–295. New York: Routledge.

Williams, Porter. 2015. "Naturalness, the Autonomy of Scales, and the 125 GeV Higgs." Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics 51:82–96.

Williams, Porter. 2021. *Renormalization Group Methods.* In *The Routledge Companion to the Philosophy of Physics*, edited by Eleanor Knox and Alastair Wilson, 296–310. New York: Routledge.