

Quantum Gravity at Low Energies

David Wallace*

December 24, 2021

Abstract

I provide a conceptually-focused presentation of ‘low-energy quantum gravity’ (LEQG), the effective quantum field theory obtained from general relativity and which provides a well-defined theory of quantum gravity at energies well below the Planck scale. I emphasize the extent to which some such theory is required by the abundant observational evidence in astrophysics and cosmology for situations which require a simultaneous treatment of quantum-mechanical and gravitational effects, *contra* the often-heard claim that all observed phenomena can be accounted for either by classical gravity or by non-gravitational quantum mechanics, and I give a detailed account of the way in which a treatment of the theory as fluctuations on a classical background emerges as an approximation to the underlying theory rather than being put in by hand. I discuss the search for a Planck-scale quantum-gravity theory from the perspective of LEQG and give an introduction to the Cosmological Constant problem as it arises within LEQG.

1 Introduction

What is a theory of ‘quantum gravity’? It is widely described¹ as any physical theory which combines general relativity and quantum theory; or, alternatively, as any physical theory that has classical general relativity and quantum theory as limiting cases or partial approximations.

And here are some common claims about a quantum theory of gravity, understood in this sense:

1. We don’t have one.²

*Department of Philosophy / Department of History and Philosophy of Science; email david.wallace@pitt.edu

¹See, e.g., Butterfield and Isham (2001, p.33): “by ‘quantum gravity’ we mean any approach to the problem of combining (or in some way ‘reconciling’) quantum theory with general relativity”; Rovelli (2008): “Quantum Gravity is the name given to any theory that describes gravity in the regimes where quantum effects cannot be disregarded”; Rickles (2008, p.262): “Quantum gravity involves the unification of the principles of quantum theory and general relativity”.

² “[t]here is still no satisfactory theory” (Butterfield and Isham 2001, p.33); “an established theory of Quantum Gravity does not yet exist” (Rickles 2008, p.265).

2. More specifically, we don't have an acceptable quantum theory of gravity in the sense that we have an acceptable quantum theory of the strong, weak or electromagnetic interactions.³
3. *Empirically* we don't really need one, because — outside exotica such as Hawking radiation, the interior of black holes, and the very early Universe — all known empirical phenomena are describable either within quantum mechanics (more specifically: within the Standard Model of particle physics) or within classical General Relativity. So any case for a quantum theory of gravity will have to be based largely on theoretical considerations.⁴

All of these claims are false. If 'quantum gravity' indeed just means a theory that combines quantum mechanics with general relativity, then we have a fairly satisfactory theory of quantum gravity; indeed, we have a theory of quantum gravity that's acceptable in pretty much exactly the sense that we have an acceptable quantum theory of electrodynamics or of the strong or weak interactions. And there is abundant observational evidence for this theory, because there are an abundance of observations that cannot be explained in their entirety either by classical general relativity or by non-gravitational quantum mechanics, but which this theory describes admirably. Indeed, an important approximation regime for this theory — often called 'semiclassical gravity' — is the (admittedly often tacit) foundation of most of contemporary astrophysics and cosmology. Other regimes are less directly tested but are still have applications in physical situations much more amenable to direct observation than the extremes usually associated with evidence for quantum gravity.

This theory is simply the quantum field theory (QFT) obtained when modern quantum-field-theoretic methods (specifically, path-integral formulations and the effective-field-theory approach) are applied to classical general relativity. It has its share of conceptual puzzles and problems, to be sure — most seriously, it ceases to be well-defined when used to describe phenomena at extremely short distances, extremely high energies, and extremely high curvatures (more precisely, at lengths which are not large compared to the Planck length, $l_P = \sqrt{\hbar G/c^3} \simeq 10^{-35}\text{m}$, and similarly for other magnitudes). But this deficiency, far from being unique to the theory, is actually shared by the Standard Model and (arguably) by classical general relativity: both predict the appearances of extremes in which they can be expected to break down.

In this paper I call it *low-energy quantum gravity*, or LEQG for short. (Ex-tant names are either too clunky ('effective-field-theory-general-relativity') or too ambiguous for my purposes ('general relativity', 'quantum general relativity')). It is worth emphasizing that 'low energy' here means 'low energy relative

³"[C]urrently there is no quantum theory of gravity in the sense that there is, say, a quantum theory of gauge fields" (Callender and Huggett 2001, p.3).

⁴"there are currently no observations demanding a quantum gravitational theory" (Callender and Huggett 2001, p.3); "there are no phenomena that can be identified unequivocally as the result of an interplay between general relativity and quantum theory" (Butterfield and Isham 2001, p.36); "If there is no empirical anomaly, then why bother with quantum gravity: why do we need such a theory?" (Rickles 2008, p.268).

to Planck scales’, which is an undemanding constraint: the heart of a supernova, and the quark-gluon plasma of the early Universe, are both low-energy regimes. The purpose of this paper is to give a presentation of LEQG and of the observational evidence that supports it, and to consider how the broader project of quantum gravity should be understood in the light of LEQG. I intend the paper as a conceptual presentation complementing more calculationally-oriented discussions, and for the most part assume without proof (but with references) the mathematical results I use. (The main exception is the somewhat-explicit presentation I give of the background-field method in sections 7–8.)

The paper is structured as follows. In sections 2–3 I review a range of astrophysical and cosmological applications which require simultaneous consideration of quantum theory and gravity for their analysis. In section 4 I present ‘semiclassical gravity’, a hybrid theory where the metric field couples to the expected value of the stress-energy tensor, as a unifying framework for these applications, and argue that the observational evidence for ‘classical’ general relativity mostly supports either semiclassical gravity or phenomenological applications of general relativity. In sections 5–6 I discuss the limitations of semiclassical gravity and the observational evidence — both at laboratory scales and in cosmology — that points beyond it, to a theory of quantized metric fluctuations.

In sections 7–8 I present LEQG in general, stressing its non-perturbative and background-independent character and explaining how background-dependent descriptions arise from it as contingent approximations. In section 9 I consider to what extent LEQG can be seen as empirically confirmed, and identify those regimes in which it is and is not supported by evidence. In section 10 I consider the breakdown of LEQG at Planckian scales and the question of how to seek a theory applicable at those scales. And in section 11 I give an introduction to the notorious *cosmological constant problem*, which arises in LEQG simply as a matter of calculation. Section 12 is the conclusion.

A note on level and rigor: I assume some background knowledge of general relativity and of quantum field theory in the path-integral formulation. I have aimed for about the level of mathematic rigor found in theoretical high-energy physics, which of course is lower than that found in pure mathematics, in some areas of general relativity, and in some parts of the foundational literature. For those who find results stated at this level unconvincing, I hope the paper will at least be illuminating as to how quantum gravity looks from the perspective of mainstream theoretical physics.

A note on sources: the particular presentation of LEQG I give here is novel in the minimal sense that I worked it out for my pedagogical purposes and have not straightforwardly followed any single source, but I claim no originality in the broader sense. Good general sources for effective-field-theory discussions of gravity, from a perspective somewhat complementary to my own and providing more on the technical details, are (Burgess 2004, Donoghue 1997, 2012; Kiefer 2007,ch.2).

Finally, a note on units: I take $\hbar = c = 1$ except where explicitly noted otherwise, so that mass has the dimensions of inverse length, and so I talk interchangeably about physics at high energies and at short lengthscales. I do

not take $G = 1$ (in low-energy quantum gravity this is often unhelpful). Note that in these units G is simply related to the Planck length, $G = \sqrt{l_P}$.

2 Observing quantum gravity: the Newtonian regime

Take a rock, or any other solid object. Hold it in your hand, a few feet above the floor, and let go. Observe its steady downward acceleration followed by a sharp deceleration when it hits the floor.

Congratulations! You just carried out an experiment in quantum gravity. (Who said fundamental physics experiments had to be expensive?) The rock accelerated downwards because of gravity: if you ignore gravity and just treat the reference frame of the room you're standing in as inertial, you won't predict that acceleration. It stopped because it and the floor are made of non-interpenetrable solid matter, and to explain *that* requires the Pauli exclusion principle (there is no classical microscopic theory of stable solid matter).

That experiment only required Newtonian gravity and quantum theory. To find an experiment that needs both *general relativity* and quantum theory, just check the map app on your cellphone. Famously,⁵ the GPS network requires general relativity for a full analysis; the processor in your phone needs quantum theory for a full analysis; the covariation of your physical position with your apparent position on the map cannot be fully analysed without both general relativity and quantum theory. Nor can the various classic experiments that use atomic clocks to measure general-relativistic time dilation.

This probably seems like cheating. In each case, there is a *background* classical gravitational field against which quantum phenomena played out. One might imagine a division of labor: classical general relativity tells us the background gravitational field (or the spacetime metric, if you prefer) and only then do we apply quantum mechanics. But it at least makes the point that we have to be a bit careful saying what we mean when we say that “there are no phenomena that can be identified unequivocally as the result of an interplay between general relativity and quantum theory” (Butterfield and Isham 2001, p.36).

Here is an attempt to be more careful: for an experiment or observation to require a quantum-gravitational description, it must require a quantum-mechanical description of *self-gravitating* matter. None of the experiments above fit this description: in each, a quantum-mechanical system (the rock, the satellite/cellphone system, the atomic clock) is moving in a gravitational field, but is not being treated as the source of a gravitational field.

But it is also easy to find situations which are quantum-gravitational in this more careful sense. Again: there is no classical microscopic account of solid matter. So a full physical description of the structure of Earth, or any other planet,⁶ requires quantum theory. But of course it also requires gravity: a

⁵For details, see (Ashby 2003).

⁶Even gas giants have solid cores; nor is this a contingent accident: self-gravitating gas

physical analysis of the structure of a planet involves solving for its density and pressure as a function of its radius, taking as input (a) the equation of state of the matter forming the planet, which requires quantum theory to explain, and (b) Newton’s law of gravity.

One might object that geologists do not in practice use quantum theory to calculate the equation of state of the Earth’s matter: perhaps we should understand that equation of state as phenomenological rather than quantum-mechanical. Replying to that objection would take us deep into both the technical details of geology and the philosophical problems of reduction and emergence, so I will bypass it by considering the technically simpler case of stars.⁷ Stars like the Sun can be reasonably well modelled, at least in their broad-brush features, as balls of dilute gas or plasma, both of which are pretty well understood in microphysical terms (at least if we bracket philosophical problems in statistical mechanics). Now granted, the equations of state for plasmas and gases can be understood classically⁸ — but there is no classical account available of nuclear fusion, the process by which a star’s temperature is maintained against the constant loss of energy through radiation. Nor is there a classical account available of how photons diffuse from the stellar core to the star’s atmosphere and so are visible as sunlight. A proper microscopic understanding of fusion requires a quantum-mechanical account of the nucleus and of the strong interaction; a proper microscopic understanding of stellar opacity requires quantum electrodynamics. So a full account of a star requires both quantum theory and gravity.

If we consider stars smaller than the sun, quantum mechanics enters even into their equations of state. In smaller stars the core is maintained by *electron degeneracy pressure*: in semiclassical terms, the Pauli exclusion principle permits only one electron of each spin in a phase-space region of size \hbar^3 , so as the spatial density of electrons becomes larger they are forced into higher-and-higher-momentum states in order to conform to the exclusion principle. (This is well-understood physics: the semiclassical and verbal gloss I give is replaceable by sharp mathematics.) The same phenomenon occurs in white dwarf stars, the corpses of stars like the Sun which have exhausted their fusion fuel. In either case, the structure of the star is explained, and indeed calculated quantitatively and checked against observational data, using a marriage of quantum theory and Newtonian gravity.

clouds aren’t stable, because of the gravothermal catastrophe (cf (Binney and Tremaine 2008, pp.567-572) and references therein).

⁷The physics I discuss here is well-understood; (Kippenhahn *et al* 2012) is a standard reference.

⁸Or at least, let’s stipulate that they can; for some reasons for doubt, see (Wallace 2016b).

3 Observing quantum gravity: relativistic matter and strong gravity

The examples I have given so far are largely nonrelativistic: we need to consider relativistic physics to handle fusion reactions, and we need to consider electromagnetic effects to understand heat flow around stars, but the short-term (hydrostatic) descriptions of planets, main-sequence stars, and white dwarfs rely just on nonrelativistic quantum mechanics and on Newtonian gravity. But it is not hard to find examples in astrophysics where Newtonian gravity is insufficient and relativistic effects must be considered, but where quantum theory remains essential.

The first class of examples arise where it matters that Einsteinian gravity couples to the full stress-energy tensor of matter, and not just to its rest mass. For instance, in cosmology we identify ‘radiation’ and ‘baryonic matter’ components of the overall matter distribution (along with the more mysterious ‘dark matter’ and ‘dark energy’ components). As a capsule summary of their physics:

- The radiation can basically be treated as black-body; its energy density is proportional to the fourth power of its temperature.
- The baryonic matter can be treated initially as a relativistic plasma, with appropriate equation of state; in due course it cools down to a nonrelativistic gas, at which point its energy density is effectively independent of its temperature, and inversely proportional to the Universe’s effective volume.
- Initially the radiation and matter are in thermal equilibrium with one another and so at the same temperature. Eventually (specifically, when protons and electrons combine to form hydrogen, making the baryonic matter transparent) the interaction between the two drops off too rapidly to maintain equilibrium and they continue to evolve separately.
- After this stage, the temperature of the radiation is inversely proportional to the universe’s effective radius, and so the energy density of the radiation falls off more rapidly than that of the baryonic matter, and the latter dominates the former gravitationally.

This is standard, well-understood cosmology⁹ — but it requires both quantum theory and general relativity to understand. The energy density of black-body radiation is, famously, a quantum-mechanical result that cannot be reproduced classically — indeed, exactly this problem led Planck to the idea of the light quantum. Similarly, the physics of how the radiation and matter move out of equilibrium is quantum-mechanical, turning on the interaction of photons first with plasma and then with atomic hydrogen. But that energy density is not a nonrelativistic ‘mass’, and to properly couple it to gravity requires general relativity. Ordinary cosmology is an application of relativistic quantum gravity.

⁹See (Weinberg 2008) or any other standard text for details.

The second class of examples arise when the strong-gravity regime of general relativity comes to the fore. For instance, consider neutron stars, white dwarfs' more extreme cousins. In a neutron star, gravitation-induced pressure is so high that protons and electrons are induced to combine into neutrons, and the resultant neutron fluid then supports itself by neutron degeneracy pressure. The basic physics is similar to white dwarfs, but the much higher densities involved mean that a general-relativistic description of a neutron star's gravity is necessary to give an accurate analysis of its structure. The physics of neutron matter is also more complicated by far than that of the degenerate matter in white dwarfs; a full understanding remains elusive, but what understanding we have relies on a combination of finite-temperature quantum field theory¹⁰ and general relativity. (For details, see (Keller *et al* 2021) and references therein.)

For an example which combines both relativistic features of matter and the strong-gravity regime of general relativity, consider gravitational supernovae, the catastrophic explosion of (some) large stars when they exhaust their nuclear fuel. In outline, sufficiently large stars build up a core of degenerate iron (iron cannot release further energy through nuclear fusion). When this core exceeds a certain mass it becomes unstable due to relativistic effects and collapses very rapidly in on itself to form a neutron star. The remaining matter falls onto the neutron core and 'bounces' (in a fashion which remains incompletely understood), dispersing into a large cloud around the neutron star, accompanied by a catastrophic release of energy, mostly in the form of neutrinos. The core of a supernova is an extreme physical environment. As David Arnett puts it in his widely-cited monograph on supernovae and nucleosynthesis,

The problem abounds with the extreme and the exotic. Neutrinos, which eluded direct detection for decades because of the extreme weakness of their interaction, are trapped in the collapse. The collapse is so fast that a mass, more than half that of the sun, shrinks in on itself at velocities approaching the speed of light. Einstein's general relativity theory becomes necessary to describe the event, not merely a subtle correction to the gravitational theory of Newton. Perhaps nowhere else in the Universe does it matter so much whether neutrinos obey Pauli's exclusion principle. (Arnett 1996, p.381)

(For further details of the physics here, see (Arnett *ibid.*) and references therein.)

To conclude this part of the discussion: it is flatly wrong that contemporary physics provides few or no situations in which gravity and quantum theory must be considered in parallel. Astrophysics and cosmology routinely give us situations in which quantum-mechanical and self-gravitation effects must be considered in parallel, and many of these situations involve core features of relativistic gravitation and relativistic quantum theory.

¹⁰In a strict sense this is not a *relativistic* application of quantum field theory — the 'finite temperatures' in question are much lower than the neutron rest mass — but it certainly lies outside the domain of applicability of nonrelativistic particle mechanics.

4 Semiclassical gravity and the observational evidence for general relativity

If astrophysics and cosmology are routinely using both quantum-mechanics and gravitational physics in the same problems, it ought to be possible to say — at least in a rough-and-ready way — what the principles are that govern their combination. And indeed, I think it is fairly straightforward to read off from these applications how this goes. In each case, quantum mechanics is used to give us appropriate dynamical equations for bulk matter, using the usual rules of quantum statistical mechanics: work out the expected values of quantities like fusion rates, pressures, and energy densities, and then treat these as the actual values of those quantities under the assumption that quantum fluctuations in their values are negligible (and if necessary check that assumption by calculating the variance in those values using the quantum formalism, though in most cases it is obvious on physical grounds that they are negligible). Normally the quantum calculations are carried out on flat spacetime, on the assumption that the dynamically relevant scales for matter are much smaller than the relevant spacetime curvature scales, and translated to general curved spacetimes via the equivalence principle. Then the matter equations thus obtained are combined with Einstein’s field equations and solved.

This approach, where quantum-mechanical matter is analysed on a classical spacetime and the two are coupled by putting the expected value of the stress-energy tensor into the Einstein field equations —

$$G_{\mu\nu}[g] + \Lambda_0 g_{\mu\nu} = 8\pi \langle T_{\mu\nu} \rangle_{|\psi\rangle} \quad (1)$$

— is generally referred to as *semiclassical gravity*. It has been reasonably widely discussed in the quantum-gravity literature (see, e.g., Kiefer 2007, pp.15-21, Wall 2012, p.8, and references therein), but it is not often recognized in that literature that — as the examples of sections 2–3 demonstrate — it has widespread applications in extant physics, encompassing much of contemporary astrophysics and cosmology.

Indeed, we can go further: in an important sense, semiclassical gravity is more strongly supported by observational evidence than ‘classical’ general relativity. In the abstract, modelling a system via general relativity requires:

1. A class of models $\langle \mathcal{M}, g, \phi \rangle$, where \mathcal{M} is a four-dimensional manifold, g is a Lorentzian metric on that manifold, and ϕ is a placeholder for whatever mathematical objects on the manifold represent the matter fields.
2. Some kind of dynamical equations for the matter fields ϕ , expressed on the background metric g .
3. A dynamical equation of the form

$$G_{\mu\nu}[g] + \Lambda g_{\mu\nu} = 8\pi \mathcal{T}[\phi, g] \quad (2)$$

where $\mathcal{T}[\phi, g]$ is some kind of expression that determines the right-hand side of the equation in terms of the matter fields and the metric.

Together, this (at least formally) forms a closed dynamical system for g and ϕ , which we can then solve to determine which of the triples $\langle \mathcal{M}, g, \phi \rangle$ represent real dynamical possibilities.

There are essentially four concrete versions of this schema that are discussed in physics. The first is *vacuum general relativity*, where there are no matter fields at all and the right hand side of the Einstein field equation is set to zero. The second is *phenomenological general relativity*, which eschews any microphysical account of matter and uses ‘dust’, ‘perfect fluid’, ‘point particle’ and similar models. The third is semiclassical gravity, which we have already discussed. And the fourth is *classical general relativity*, where a classical matter field like the Klein-Gordon field, Dirac field or electromagnetic field is coupled to gravity, usually using an action principle to determine both the dynamical equations for matter and the stress-energy tensor.

The first three of these find fairly widespread application in astrophysics. Vacuum general relativity is applied to the study of black holes and gravity waves; the calculations underpinning the LIGO gravity wave observatory’s fit to predictions rely on it (cf. (Coleman Miller and Yunes 2019) and references therein). Phenomenological general relativity (albeit often in a Newtonian approximation) is commonly applied in various pieces of galactic-dynamics and cosmology, where the details of matter can be abstracted away, and also in solar-system physics, where Birkhoff’s theorem¹¹ allows us to disregard the details of the Sun’s structure given that it is spherically symmetric. And we have seen that semiclassical gravity is necessary whenever we need a non-phenomenological description of quantum-mechanical matter. In practice the line between phenomenological general relativity and semiclassical gravity is blurry: it is often a matter of degree in astrophysics to what extent the dynamical equations of matter are derived bottom-up and to what extent they are modelled phenomenologically.

Classical general relativity, however — despite its undoubted mathematical beauty — has no applications at all to observational data (except insofar as we treat it as including vacuum general relativity as a special case). To the best of my knowledge, all observationally-relevant treatments of matter coupled to general relativity use some mix of phenomenological and quantum-mechanical treatments: there are no observationally-relevant cases where a classical field, in the usual sense of that term, is coupled to gravity. Nor should this be surprising: the only classical fields seem to be¹² gravity itself and the electromagnetic field, with the other quantum fields of the Standard Model manifesting in the classical regime, if at all, as particles. And while in principle a classical electromagnetic field could generate a gravitational effect, in practice this does not seem to occur in astrophysically-relevant situations.¹³ (Black-body electromagnetic radiation

¹¹See, e. g. , (Wald 1984, p.125) or (Carroll 2003, p.197).

¹²There are some field-theoretic models of dark matter (see (Urena-Lopez 2019) for a review), so it would be premature to absolutely rule out a third classical field.

¹³The nearest I can imagine to an exception might be the enormously strong magnetic fields generated by neutron stars and by black hole accretion disks. But even there, a self-contained account would require a quantum-mechanical and/or phenomenological description

self-gravitates non-trivially in some circumstances, but as we have already noted, this requires a quantum-mechanical treatment.)

So: what is ‘general relativity’, if we mean the theory which is actually used and applied in observational physics? It is not a self-contained classical theory that could be thought of as fundamental were it not for the pesky quantum effects in other regimes: it is in practice a higher-level effective theory, in which matter is treated either phenomenologically or quantum-mechanically through the methods of semiclassical gravity.¹⁴

5 The limitations of semiclassical gravity

Semiclassical equations like (1) are very common in physics. Known variously as ‘Vlasov equations’, ‘mean-field approximations’, and ‘Hartree(-Fock) approximations’, they have extremely wide applicability, from nuclear and atomic physics through to galactic dynamics. The Hartree approximation to nonrelativistic quantum electrostatics, for instance —

$$i\hbar\frac{\partial\psi}{\partial t}(x,t) = -\frac{\hbar^2}{2m}\nabla^2\psi(x,t) - eV(|\psi|)(x,t)\psi(x,t) \quad (3)$$

where

$$V(|\psi|)(x,t) = \int dx' \frac{e|\psi(x',t)|^2}{|x-x'|} \quad (4)$$

— describes electrons moving in the mean electric field created by those same electrons, and is the starting point for much of atomic physics and quantum chemistry.¹⁵ But normally these equations are taken as approximations to the underlying dynamics, approximations which hold only on the assumption that multiparticle correlations, or their field-theoretic equivalents, are low. (The Hartree equation, for instance, is derivable¹⁶ as a first-order approximation to the N -particle Schrödinger equation under the approximation that the quantum state remains approximately in product form). And there is a clear physical reason why these approximations ought to fail in the presence of strong correlations: if the quantum state of N electrons is a superposition of the electrons at X and the electrons at Y , an $N + 1$ th electron interacting with the first N would be

of the generation of the field, and could not be analyzed just using the classical theory of electromagnetism coupled to gravity. A black hole with a significant electric charge *would* be modelable in that classical theory, but such objects do not appear to occur in nature, for physically obvious reasons.

¹⁴It is also non-fundamental for another reason lying somewhat outside the main scope of this paper: in at least a large fraction of applications of general relativity, including almost all applications at very large scales, the metric tensor g is not the ‘true’ metric (whatever that means) but a highly coarse-grained approximation to that metric, smoothing out metric inhomogeneities over macroscopically large regions. Cosmology often smooths out even the structure of galaxies, let alone the structure of the matter within them. See (Wiltshire 2011) and references therein for further discussion of this point.

¹⁵Strictly speaking, the equation used is the Hartree-Fock equation, which takes better account of the antisymmetrisation properties of a multi-electron wavefunction.

¹⁶See, e. g., (Bonitz 2016).

expected to end up entangled with them, not feeling the electric field as if its source was at $(X + Y)/2$. Adopting the language of decoherent histories,¹⁷ an electron in a history h ought to feel the electric field of the other electrons at their positions as described in h , not their expected position averaged over all histories.

It is then natural to expect that semiclassical gravity is likewise an approximation to some more fully quantum theory, and that the equations of semiclassical gravity would fail under the same circumstances that other semiclassical methods fail: that is, when the quantum state is a superposition of substantially different mass distributions (something that is not the case¹⁸ in any of the astrophysical and cosmological examples we have considered so far). And it is fairly straightforward to test this empirically: simply use a quantum-mechanically random process (a spin measurement, say) to place a mass large enough to generate a detectable gravitational field in one of two macroscopically distinct locations in the lab; then measure the gravitational force. If semiclassical gravity fully describes quantum gravity in this regime, the force should be detected as directed towards an empty location intermediate between the two locations; if semiclassical gravity is a mean-field approximation to something else, we would expect the force to be detected as directed towards the observed location of the mass. This experiment has actually been done (Page and Geilker 1981), demonstrating (perhaps unsurprisingly) that semiclassical gravity indeed fails in this situation.

Similarly, considerations of the long-term chaotic behavior of the solar system, in which one would expect the locations of the planets to become macroscopically superposed on megayear timescales (Zurek and Paz 1999), might be said to provide an observational version of the Page-Geilker experiment, as Page and Geilker themselves note (*ibid.*, p.980). On sufficiently long timescales, no unitarily-evolving quantum state will remain macroscopically definite.¹⁹

The Page-Geilker experiment has received some pushback in the philosophy literature, notably by Mattingly (2005, 2009). Mattingly's objection is that Page and Geilker assume the Everett interpretation of quantum mechanics (they do indeed mention it explicitly in their paper) and that their experiment does not have any implications for semiclassical gravity without this assumption.

I have no solution to the quantum measurement problem. Nor do I find wholly satisfying any of the extant proposals for solving it. Yet I cannot worry overmuch about a proposal for a gravitation theory merely because it fails to solve the problem, or rather because it

¹⁷(Gell-Mann and Hartle 1989, 1993; Griffiths 1984, 1993; Omnes 1988, 1992); see also reviews by Halliwell (1995) and Wallace (2012, ch.3).

¹⁸There is a caveat here: given unitary dynamics, the *universal* quantum state (insofar as such a thing makes sense) is certainly in a superposition of different mass distributions. I discuss this below.

¹⁹For that reason, at least on interpretations which treat the quantum state as observer-independent, the empirical evidence for semiclassical gravity really supports what we might call *branch-relative semiclassical gravity*: semiclassical gravity applied not to the universal state but to the branch-relative state.

undermines one or two of the proposed solutions. (Mattingly 2005, pp.330-331).

But this understates the scope of the Page-Geilker result (and of similar, more theoretical arguments; cf (Kiefer 2007, pp.15–22) and references therein). The argument I sketched above does not rely on any *interpretative* assumptions, but just on the *mathematical* assumption that the quantum state of the lab in which the experiment is conducted evolves unitarily (or, more precisely, according to the modified version of unitarity that incorporates the semiclassical gravitational field). The Everett interpretation makes this assumption, but it is not alone in doing so: the various consistent- and decoherent-history approaches to quantum mechanics all assume unitary dynamics, as do the various hidden-variable and modal ‘interpretations’.²⁰ Even pragmatist or instrumentalist approaches such as QBism (Fuchs *et al*, 2014), Healey’s pragmatism (Healey 2012, 2017), Fuchs’ and Peres’ operationalism (Fuchs and Peres 2000), or most versions of the Copenhagen interpretation at least permit the lab to be analyzed unitarily, since they all permit Wigner’s-friend scenarios where an external observer considers the lab (experimenters and all) as a single quantum system. This same move is permitted on at least some interpretations of Dirac and von Neumann’s ‘measurement-collapse’ approach, since that approach is often taken to allow us to put the transition between unitary and non-unitary dynamics (von Neumann’s ‘cut’) as close to the observer as we wish. It is permitted, too, under Rovelli’s relationalist interpretation (Rovelli 1996). Of the various interpretations of quantum mechanics, and proposals for modifications of quantum mechanics, currently extant in the philosophy and physics literature, only dynamical-collapse theories clearly prohibit a unitary analysis of the Page-Geilker experiment. And while the suggestion that dynamical collapse might have some connection to gravity is certainly intriguing, and has been the source of interesting theoretical speculations and experimental proposals (to which I return in section 9), at this point we are well outside anything that might be called an ‘interpretation’ of quantum mechanics. The mainstream use of quantum theory in modern physics (cf Wallace 2019c, 2016a) assumes unitary quantum dynamics, and that is all that is required for Page-Geilker-type experiments to falsify semiclassical gravity.

Nonetheless, the Page-Geilker experiment — and Zurek and Paz’s theoretical models — have a pretty ‘foundational’ feel to them, concerning as they do macroscopic superpositions which are implied by the equations but not in any direct sense observable. There is, however, an important class of applications of quantum gravity in mainstream cosmology where the presence of quantum

²⁰Hidden-variable theories, to be sure, also allow for the possibility that the gravitational field is sensitive to the hidden variables and not (or not just) to the physical features encoded in the quantum state; see (Pinto-Neto and Struyve 2018, section 8) for further development of this idea. However, theories of this kind take us beyond semiclassical gravity, in which the right-hand-side of the Einstein equation is a quantum expectation, independent of the hidden variables. (For that reason they also face a potentially severe observational challenge in relativistic regimes, where a quantum calculation of the stress-energy tensor is used in astrophysics and cosmology.)

fluctuations in the gravitational field, and the responsiveness of matter to those fluctuations and not just to their expected value, has direct observational consequences. These applications — to the quantum origins of inhomogeneity in the microwave background radiation and in the cosmological-scale distribution of matter — will be of interest both because they further mark the limitations of semiclassical gravity and because they provide further examples of the direct use of quantum gravity in contemporary physics.

6 Quantum fluctuations in the inflaton field

A key feature of gravitational physics, owing its origin to the universally-attractive nature of gravity, is that small fluctuations from isotropy and homogeneity in the matter distribution will grow exponentially in size over time (rather than be damped back to uniformity as in the case of, say, a non-self-gravitating dilute gas or electromagnetic plasma). As such, even very small fluctuations away from uniformity in the early Universe can suffice to generate the highly non-uniform cosmos we observe today, in which typical matter densities in stars exceed the average matter density of the Universe by around thirty orders of magnitude. When in 1992 the COBE satellite detected relative fluctuations of 10^{-5} in the observed cosmological microwave background (CMB), this began the modern era of understanding in quantitative detail exactly how these small fluctuations indeed seed the large-scale structure we observe today, and arguably can be seen as the beginning of the current golden age of observational cosmology. (Here and for the rest of this section, see (Weinberg 2008) for details and references.)

The *origin* of these fluctuations is another matter. As long as we assume a classical, deterministic picture of cosmology, small fluctuations in the CMB must have their origin in still-smaller — but still non-zero — fluctuations in the still-earlier Universe. But the dominant view in modern cosmology is that ultimately we should *not* assume that ‘classical, deterministic picture’, and instead should look to understand these primordial fluctuations as the classical remnants of *quantum fluctuations* in the early universe, even given an isotropic and homogeneous quantum state for the early universe — and that concrete quantum theories of those fluctuations can be used to calculate a predicted spectrum of fluctuations, compared to observation, and thus tested. The dominant such theory, *inflation*, has indeed been so tested, and while the ‘test’ in question is scarcely precise, inflation does indeed correctly predict the form of the CMB spectrum. This connection of inflation to the CMB and thus, indirectly, to large-scale structure formation has probably now replaced older, more foundational, motivations as the main reason for its popularity amongst cosmologists (see (Smeenk 2017) for more discussion).

It should be clear on physical grounds that any such quantum analysis of fluctuations cannot assume semiclassical gravity. If we begin with a homogeneous and isotropic quantum state describing matter on a homogeneous and isotropic background, and allow it to evolve semiclassically, its expectation values will remain homogeneous and isotropic. From the quantum point of view,

structure formation is a correlational phenomenon invisible to mean-field methods. And indeed, direct examination of the sorts of calculations used in this field confirms that they are not made within semiclassical gravity. Instead, the usual approach is to split the metric field g into a sum $g = g_0 + h$, where g_0 is homogeneous and isotropic and h thus represents fluctuations around g_0 , and then to quantize h and the matter field (say, an inflaton field ϕ) as quantum fields interacting with one another on a background spacetime with metric g_0 . Fluctuation information is then contained in expectation values like $\langle \hat{h}(x)\hat{h}(y) \rangle$, which, given decoherence, can be interpreted as encoding correlations between fluctuations at different points.

In the most straightforward of these analyses, the radiational degrees of freedom of h (i.e., the gravitons) are assumed neglectable, and so h is basically a functional of the matter field, representing gravitational interactions between matter fluctuations. In more detailed analyses (e.g., Bhattacharya *et al*, 2006) gravitons are also present. But even when gravitons are frozen out, this is fairly clearly a fully quantum theory of gravity, with the gravitational field being a superposition of different classical values, and entangled with the matter field.

It is fair to say that the physics here is significantly more speculative and significantly less observationally constrained than the very mature astrophysics and cosmology that underlies applications of semiclassical gravity. But it is also fair to say that it is much *less* speculative, and much *more* observationally constrained, than proposals for Planck-scale quantum gravity. Inflationary models do make predictions about the quantitative form of the CMB fluctuations; those predictions are non-trivial; they have so far been confirmed. And active efforts continue to further constrain inflationary models with observational data: notably, the BICEP and Keck arrays at the South Pole²¹ aim to measure polarization in the CMB, which is an indirect measure of the graviton background in the primordial Universe.

7 Background field methods

The applications I have so far discussed demonstrate fairly widespread applications of quantum gravity in contemporary astrophysics and cosmology, and comparably widespread exposure of those applications to observational test. But at first sight they also paint a rather messy disunified picture of quantum gravity at low energies. Semiclassical quantum gravity has a fully dynamical interaction between quantum matter and spacetime metric, but treats the latter as fully classical, and is for that reason (as we have seen) not satisfactory except as some sort of approximation. But it is not obviously any kind of approximation to the essentially perturbative approach to quantum gravity used in inflation, in which a classical background spacetime plays a central role and in which quantization is apparently only defined relative to that background.

This perception is misleading. The methods of modern quantum field theory — specifically, the ‘background field method’, developed by Goldstone, Salam

²¹See Ada *et al* (2018) and references therein.

and Weinberg in particle physics, and independently by de Witt in the context of quantum gravity (see Weinberg (1995, ch.16) for a review and for original sources) provide an elegant way to see both semiclassical methods and quantize-on-a-background methods as different approximation schemes for the same underlying theory. In this section I give a brief outline of the background field method in the simpler context of scalar field theory; section 8 applies the method to gravity.²²

In contemporary QFT, the theory is normally specified by means of a *path integral*, which for scalar field theory has the form

$$\langle \phi_f; t_f | \phi_i; t_i \rangle = \int_{\phi(t_i)=\phi_i}^{\phi(t_f)=\phi_f} \mathcal{D}\phi e^{iS[\phi]} \quad (5)$$

where

- $|\phi_i; t_i\rangle$ is a quantum state in the Heisenberg representation of the theory, defined by $\widehat{\phi}(\mathbf{x}, t_i) |\phi_i; t_i\rangle = \phi_i(\mathbf{x}, t_i) |\phi_i; t_i\rangle$, and similarly for $|\phi_f; t_f\rangle$.
- $S[\phi]$, the *action* for the scalar field theory, is a quasi-local functional of the classical scalar field ϕ , which we assume takes the general form

$$S[\phi] = \int dx^4 \left(\frac{1}{2} \partial^\mu \phi(x) \partial_\mu \phi(x) - V(\phi(x)) \right). \quad (6)$$

That is, the amplitude for transitions between configuration ϕ_i at time t_i , and configuration ϕ_f at time t_f , is expressed as a path integral over all histories that connect ϕ_i to ϕ_f , with each history weighted by a phase factor determined by the action of the history. As stated this is all rather formal: the path integrals here are divergent and the theory must be regularized through the addition of some high-energy cutoff in order for the path integrals to be well-defined; I postpone consideration of these issues to section 11.

Suppose we now fix a quantum state $|\text{in}\rangle$, which we choose to encode a known state of the system at some early time. The physical features of this system at all times are encoded in expectation values like

$$\langle F[\widehat{\phi}] \rangle_{\text{in}} \equiv \langle \text{in} | \widehat{\phi}(x_1) \cdots \widehat{\phi}(x_n) | \text{in} \rangle \quad (7)$$

and these can be expressed in terms of an appropriately manipulated path-integral: schematically, as

$$\langle F[\widehat{\phi}] \rangle_{\text{in}} \propto \int_{\text{in}} \mathcal{D}\phi F[\phi] e^{iS[\phi]} \quad (8)$$

where ‘*in*’ schematically indicates the boundary conditions on the integral required to enforce $|\text{in}\rangle$ as our choice of quantum state. The most common choices

²²For calculational uses it is more convenient to develop the background field method using the powerful and elegant *quantum effective action* (see (Peskin and Schroeder 1995, ch.11) for an introduction and (e.g.) Weinberg (1995, ch.16) or Banks (2008, ch.2) for more advanced treatments). But this is overkill for my purposes; here I adopt a more heuristic approach.

of state are the vacuum state $|\Omega\rangle$ and (through a generalisation to mixed states) the thermal state $\widehat{\rho}_\beta = \exp(-\beta\widehat{H})/\text{Tr}\exp(-\beta\widehat{H})$ which describes a system at equilibrium at temperature $T = 1/k_B\beta$, but we can in principle choose any state we wish, including states which are not eigenstates of energy (or mixtures of eigenstates) and so describe time-dependent phenomena. (In this latter case, the integral is the so-called ‘in-in’ or ‘Schwinger-Keldysh’ path integral, where the integral is taken on a double contour in time that extends into the far future and then back again, effectively doubling the fields present; for the details, see (e.g.) Berges (2004, section 3.1) or Calzetta and Hu (2008, ch.6).) From here on, I will usually omit the ‘in’ subscripts, taking the specification of the state as tacit.

If the expectation value $\phi_0(x) = \langle \text{in} | \widehat{\phi}(x) | \text{in} \rangle$ is zero, and if a number of other technical conditions hold, then we can evaluate these path integrals perturbatively and get increasingly good approximations for them. In these cases, we can normally (at least approximately, at least on short timescales) interpret the QFT in the standard way as a theory of spin-0 particles, and derive the masses, decay rates, and scattering cross-sections for those particles from the path integral. However, in many cases of interest ϕ_0 is not zero, and these perturbative methods cannot straightforwardly be applied; this is the context for the background field method.

To begin with, suppose we carry out an infinitesimal change of variables in the path integral: $\phi(x) \rightarrow \phi(x) + \epsilon(x)$. The value of an integral is invariant under any such change, so

$$\begin{aligned} 0 &= \int \mathcal{D}\phi e^{iS[\phi+\epsilon]} - \int \mathcal{D}\phi e^{iS[\phi]} = \int \mathcal{D}\phi \int dx^4 i \frac{\delta S}{\delta \phi(x)} \epsilon(x) e^{iS[\phi]} \\ &= i \int dx^4 \epsilon(x) \int \mathcal{D}\phi \frac{\delta S}{\delta \phi(x)} e^{iS[\phi]} \\ &\propto i \int dx^4 \epsilon(x) \left\langle \frac{\delta S}{\delta \widehat{\phi}(x)} \right\rangle \end{aligned} \quad (9)$$

so that

$$\left\langle \frac{\delta S}{\delta \widehat{\phi}(x)} \right\rangle = 0. \quad (10)$$

(Since the variational derivative of S is just the Euler-Lagrange equation, i. e. the classical field equation for the action S , this can be thought of as the path-integral version of Ehrenfest’s theorem.)

Now we make a *finite* change of variables,

$$\chi(x) = \phi(x) - \phi_0(x). \quad (11)$$

(Recall that ϕ_0 is the expectation value of the field on our specified quantum state $|\text{in}\rangle$.) We can now rewrite the action as

$$S[\phi_0 + \chi] = S[\phi_0] + \left. \frac{\delta S}{\delta \phi} \right|_{\phi_0} \cdot \chi + \Sigma_{\phi_0}[\chi] \quad (12)$$

where in the second term on the right I use the abbreviated notation

$$\left. \frac{\delta S}{\delta \phi} \right|_{\phi_0} \cdot \chi \equiv \int dx^4 \left. \frac{\delta S}{\delta \phi(x)} \right|_{\phi_0} \chi(x). \quad (13)$$

Plugging this into (10), we obtain

$$0 = \left. \frac{\delta S}{\delta \phi} \right|_{\phi_0} + \left. \frac{\delta}{\delta \phi} \right|_{\phi_0} \cdot \left. \frac{\delta S}{\delta \phi} \right|_{\phi_0} \langle \chi(x) \rangle + \left\langle \left. \frac{\delta \Sigma_\phi}{\delta \phi} \right|_{\phi_0} \right\rangle. \quad (14)$$

The second of these terms vanishes exactly, because by construction the expectation value of $\hat{\chi}$ is zero, so that this simplifies to

$$0 = \left. \frac{\delta S}{\delta \phi} \right|_{\phi_0} + \left\langle \left. \frac{\delta \Sigma_\phi}{\delta \phi} \right|_{\phi_0} \right\rangle. \quad (15)$$

This is the sum of the classical Euler-Lagrange equation for the scalar field, and a quantum expectation value that can be expressed as a path-integral,

$$\left\langle \left. \frac{\delta \Sigma_\phi}{\delta \phi} \right|_{\phi_0} \right\rangle \propto \int \mathcal{D}\chi \left. \frac{\delta \Sigma_\phi[\chi]}{\delta \phi} \right|_{\phi_0} \exp i \left(\chi \cdot \Sigma_{\phi_0}[\chi] + \left. \frac{\delta S}{\delta \phi} \right|_{\phi_0} \right) \quad (16)$$

where I have just used the change of variables (11) in (8) and absorbed a constant phase factor into the constant of proportionality.

We have now reexpressed the original quantum field theory as a theory of two fields: a classical scalar field, and a quantum field whose action

$$S'_{\phi_0}[\chi] = \Sigma_{\phi_0}[\chi] + \chi \cdot \left. \frac{\delta S}{\delta \phi} \right|_{\phi_0} \quad (17)$$

depends on that classical field. (The second term in this expression can be thought of as encoding quantum corrections to the classical expectation value and can largely be ignored in calculations, serving only to cancel certain diagrams in a perturbative expansion; it is the first term Σ_{ϕ_0} that really defines the interesting features of the quantum field.) The dynamics of the two fields have the classic form of a self-consistency problem: the quantum field evolves against a background defined by the classical field; the classical field responds to the quantum field via equation (15). The expectation value of this new quantum field is zero by construction, and so we can hope (though must check explicitly) that we have a well-behaved perturbative description of that field.

The best-known uses of the background field method are in spontaneous symmetry breaking (SSB), where (for the real field we are presently considering) we begin with some potential like

$$V(\phi) = -\frac{1}{2}\mu^2\phi^2 + \frac{\lambda}{4!}\phi^4 \quad (18)$$

for the action (6). V has a maximum at $\phi = 0$ and minima at $\phi = \pm(6/\lambda)^{1/2}\mu$; in elementary treatments of SSB we pick one of these minima and quantize around it (much as we considered quantizing around a particular solution of the Einstein equations in section 6). But the background-field method lets us give a fully quantum treatment, where we extremize not the classical action but the quantum effective action. The latter is dominated by the classical action but has quantum corrections which can be thought of as polarisations of the vacuum. These do not change the qualitative form of the function²³, so that its extrema are homogeneous fields equal to 0 (a maximum) or $\pm\alpha$ for some nonzero α (a minimum). Calculating the two-point functions for the quantum fields defined around these extrema reveals that the maximum corresponds to a highly unstable quantum state and so does not describe a genuine vacuum, but that the minima are the expectation values of well-behaved vacuum states. The existence of multiple extrema tells us that the theory has multiple vacuum sectors, with the actual vacuum fixed mathematically by the boundary conditions, physically by contingent facts about the field-theoretic system being studied. And crucially, we have obtained these sectors not by applying different quantization procedures for fluctuations around a background, but by defining the quantum theory once and for all by the path integral (5) and then exploring its space of solutions.

If we analyze SSB for a finite-temperature state, the expectation value will now contain thermal contributions as well as vacuum polarization. These generally have the effect of decreasing the expectation value ϕ_0 , until for sufficiently high temperature Γ ceases to have symmetry-breaking minima and the symmetry is restored (in particle-physics contexts, this is expected to happen in the early universe).

A more complicated application of the background-field method is to situations where the quantum system starts in a strongly time-dependent state (i.e., not a vacuum state, a thermal state, or a state describable as a small perturbation of either). As already noted, this requires the more-complicated ‘in-in’ path integral, which somewhat complicates the technical details of the background-field formalism. But the basic structure remains: we can use equation (15) to express the dynamics of the system as a coupled interaction of a classical field and the quantized fluctuations around that field. The underlying theory remains fully quantum-mechanical, and this particular decomposition into background and fluctuation is solution-relative.

As a final observation before we turn to gravity: *formally* the semiclassical equation (15) is exact. But it gets its practical power from the ability to calculate a good approximation to the expectation-value term, and this in turn requires (among other things) that the strength of the fluctuations (expressed by the two-point function $\langle\hat{\chi}(x)\hat{\chi}(y)\rangle$) are not too large compared to other relevant scales (the so-called ‘mean-field’ regime). If this assumption is violated, equation (15) becomes intractable, and physically the system may be better described as a

²³This need not be the case. There are well-known examples (notably the Coleman-Weinberg model (Coleman and Weinberg 1973) where the classical action has a minimum with unbroken symmetry and the symmetry breaking has an entirely quantum origin.

superposition of macroscopically different backgrounds with fluctuations around each. In the nonequilibrium case this can happen dynamically if initially small fluctuations can grow over time, causing a breakdown over time of the mean-field assumption. Exploring the physics here beyond the mean-field regime is the domain of the comparatively new field of *non-equilibrium quantum field theory* (see, e. g., (Berges 2004; Calzetta and Hu 2008) for details).

8 Low-energy quantum gravity

Formally speaking, we can define a quantum theory of gravity using path integrals, just as we defined a scalar field theory above. We want a theory of the gravitational metric g and some matter fields which we write schematically as ϕ (read ϕ as a scalar field if you like, though I intend it as a more general placeholder), and we assume an action of the form

$$S[g, \phi] = \frac{1}{8\pi G} S^V[g] + S^M[\phi, g]. \quad (19)$$

Here $S^V[g]$, the *vacuum action*, has form

$$S^V[g] = \int dx^4 (R(x) + \Lambda_0) \sqrt{-g(x)} + \text{higher-order terms in } g \quad (20)$$

— that is, it is the sum of the standard Einstein-Hilbert action, a cosmological-constant term, and unspecified higher-order terms in g . S^M is a matter action, depending on both matter and metric: for a real scalar field it would have form²⁴

$$S^M[\phi, g] = \int dx^4 \left(\frac{1}{2} g^{\mu\nu}(x) \partial_\mu \phi(x) \partial_\nu \phi(x) - V(\phi(x)) \right) \sqrt{-g(x)}. \quad (21)$$

The field theory is then specified by the path integral

$$\langle h_f, \phi_f | h_i, \phi_i \rangle = \int_{(g, \phi)|_{\mathcal{S}_i}=(h_i, \phi_i)}^{(g, \phi)|_{\mathcal{S}_f}=(h_f, \phi_f)} \mathcal{D}g \mathcal{D}\phi e^{iS[g, \phi]} \quad (22)$$

where the boundary conditions indicate that we are integrating over fields between some initial and final 3-surfaces \mathcal{S}_i , \mathcal{S}_f , and that the fields on \mathcal{S}_i must equal (h_i, ϕ_i) and similarly for \mathcal{S}_f .

This simple formal expression hides a host of conceptual and technical complications, of course, among them:

1. The path integral again requires regularization to be well-defined, and that regularization becomes more complex technically and conceptually due to the non-renormalizability of the gravitational action (see section 11 for further discussion of this point).

²⁴Arguably a non-minimally-coupled term $R\phi\sqrt{-g}$ should also be included: it violates no symmetry, and will be reintroduced by renormalization if not added initially.

2. The diffeomorphism invariance of S means that a naive attempt to calculate the path-integral will overcount and integrate over fields that are related by a (small) diffeomorphism. That is conceptually undesirable (small diffeomorphisms are normally interpreted as purely redescriptive²⁵) and leads to a pathological theory. The formal path integral $\mathcal{D}g\mathcal{D}\phi$ must be defined so as to integrate over only one element of each gauge orbit, which in actual calculations normally requires Faddeev-Popov ghosts, BRST methods, or other such heavy machinery (see (Weinberg 1995, ch.15-17) and references therein).
3. Strictly speaking, the action must be supplemented by a boundary term to get a well-defined theory, due to the presence of second derivatives of the metric in the Einstein-Hilbert action; this has no classical consequences but can lead to non-perturbative quantum effects (see (Hawking 1979, pp.749-751) or (Padmanabhan 2014) for a discussion).
4. On a more conceptual level, diffeomorphism invariance leads to the notorious 'problem of time' (see (Butterfield and Isham 1999; Rickles 2008) for discussion and further references), which can make it difficult even to interpret the formalism (at least in the cosmological case; for applications of quantum gravity to isolated systems like stars or black holes, asymptotically flat boundary conditions simplify matters). This shows up in the path integral in the fact that —unlike in the scalar case — I have not explicitly specified initial and final times. In isolated-system applications, we do actually need to specify those conditions, since the asymptotic boundary conditions make them well-defined; in the cosmological case, things are more complicated.

But notwithstanding these subtleties and open questions, the theory is sufficiently well-defined that many calculations can be done in it, and we will see that it offers a unified starting point from which can be derived the various approximations discussed previously. In particular, we can apply the background-field method of the previous section, to express the metric as the sum of its expectation value g_0 and a new field h which tracks fluctuations around that expectation value:

$$g(x) = g_0(x) + \sqrt{8\pi G}h(x) \quad (23)$$

(the reason for the $\sqrt{8\pi G}$ factor will be apparent shortly). Expanding the action around g_0 in this way, we get

$$S[g_0 + \sqrt{8\pi G}h, \phi] \equiv \frac{1}{8\pi G}S^V[g_0] + \frac{1}{\sqrt{8\pi G}} \left. \frac{\delta S^V}{\delta g^{\mu\nu}} \right|_{g_0} \cdot h^{\mu\nu} + S^M[g_0, \phi] + \Sigma_{g_0}[h, \phi] \quad (24)$$

²⁵High-energy physics does not treat *large* diffeomorphisms — those that are non-vanishing at infinity — as redescriptive, and there is no requirement to avoid integrating over fields related by large diffeomorphisms; see Belot (2018), Wallace (2019a) for discussion of the conceptual issues here.

and can then obtain a semiclassical analogue of (15) for g_0 :

$$0 = \frac{1}{8\pi G} \left. \frac{\delta S^V}{\delta g(x)} \right|_{g_0(x)} + \left\langle \frac{\delta(S^M[g_0, \cdot] + \Sigma_{g_0})}{\delta g_0(x)} \right\rangle \quad (25)$$

where the expectation value on the right hand side is calculated for a field theory with fields h, ϕ , defined on a spacetime background g_0 and with action

$$S'_{g_0}[h, \phi] = S^M[g_0, \phi] + \Sigma_{g_0}[h] + (\text{linear term in } h) \quad (26)$$

where, as before, the linear term can mostly be neglected. Since the stress-energy tensor in classical general relativity is defined as the variational derivative of the action with respect to the metric, we can interpret the second term in (25) as the expectation value of the stress-energy tensor for the h and ϕ fields. Since the first, classical, term is the sum of the Einstein tensor, the cosmological-constant term, and higher-derivative terms, we can recognize (25) as a form of the semiclassical Einstein field equation.

It is instructive to explicitly expand $S[g_0 + h, \phi]$ in powers of $\sqrt{8\pi G}$.²⁶ We obtain

$$\begin{aligned} S[g_0 + \sqrt{8\pi G}h, \phi] &= \frac{1}{8\pi G} S^V[g_0] \\ &+ \frac{1}{\sqrt{8\pi G}} \left. \frac{\delta S^V}{\delta g^{\mu\nu}} \right|_{g_0} \cdot h^{\mu\nu} \\ &+ (S^M[g_0, \phi] + S_{g_0}^{GRAV}[h]) \\ &+ \sqrt{8\pi G} \left(T_{\mu\nu, g_0}^M[\phi] h^{\mu\nu} + \frac{1}{3} T_{\mu\nu, g_0}^{GRAV}[h] h^{\mu\nu} \right) \\ &+ O(8\pi G). \end{aligned} \quad (27)$$

The first term here is the classical vacuum action, evaluated for the background field. Passing over the second term for the moment, the third term is the sum of the classical action defined on the background g_0 , and the (free) *graviton action* on that background:

$$\begin{aligned} S_{g_0}^{GRAV}[h] &\equiv \frac{1}{2} \left. \frac{\delta^2 S^V}{\delta g_c^2} \right|_{g_0} (h, h) \\ &= \frac{1}{4} \int dx^4 \sqrt{-g_0} \nabla_\alpha h_{\mu\nu} \nabla_\beta h_{\rho\sigma} [g_0^{\alpha\beta} g_0^{\mu\nu} g_0^{\rho\sigma} + 2g_0^{\alpha\sigma} g_0^{\mu\rho} g_0^{\nu\sigma} - (\nu \leftrightarrow \rho)] \\ &+ \text{cosmological-constant term} \\ &+ \text{higher terms.} \end{aligned} \quad (28)$$

²⁶The attentive reader will observe that G is dimensionful, which calls into doubt the interpretation of any such expansion. A more careful treatment (which I assume tacitly) would expand in powers of $(\sqrt{8\pi G}/L)$, where L is the typical scale of the physical processes being analyzed (e.g., the inverse momenta of the particles). Since in our units \sqrt{G} is the Planck length, this will be extremely small for low-energy phenomena and defines a meaningful expansion. The failure of this process when L approaches Planck scales points to the nonrenormalizability of the theory and its consequent failure at those scales, discussed further in section 10.

Physical fields in QFT are normalized so as to avoid scale factors in front of the kinetic term for those fields, so this term is the reason for the $\sqrt{8\pi G}$ term in the definition of h .

Returning to the second term, we can now interpret it as a source term for the creation of single gravitons. Since such a process would shift the expected value of \hat{h} away from zero, we can now see that the effects of this term will be exactly cancelled by other terms, and indeed this can be easily shown. At the technical level, the term just serves to cancel so-called ‘tadpole diagrams’ in the evaluation of the path integral; interpretatively, it can be seen as compensating for quantum corrections to $\langle h \rangle$.

The fourth term is the lowest-order coupling of gravitons to matter and to each other. In each case, the gravitons couple linearly to the stress-energy tensor of the respective excitations: T^M is defined in the standard way as the variational derivative of the matter action with respect to the (background) metric, and T^{GRAV} is likewise the variational derivative of the free-graviton action. (The latter term is actually nonlinear because T^{GRAV} itself depends on h , which also accounts for the 1/3 term in that term: the variational derivative of $T^{GRAV}[h]_{\mu\nu}h^{\mu\nu}$ is $3T^{GRAV}[h]$, not $T^{GRAV}[h]$.) Higher-order terms encode higher-order couplings.

If we now evaluate the effective action to lowest order, dropping all positive powers of $8\pi G$, the semiclassical equation for g_0 becomes

$$0 = \frac{1}{8\pi G} \left. \frac{\delta S^V}{\delta g_c(x)} \right|_{g_0(x)} + \langle (T^M + T^{GRAV}) \rangle \quad (29)$$

where the expectation value is calculated for a QFT consisting of the matter field ϕ on the background g_0 along with a field of free gravitons on that same background. We can recognize this as semiclassical gravity: a theory of quantum matter evolving on a classical spacetime background, with the backreaction of the quantum matter on the background given by Einstein’s equation with the expected value of the stress-energy tensor as source term. (The graviton term is often omitted from presentations of semiclassical gravity, but — at least in high-energy physics — it is generally understood to be present in a careful formulation of the theory, and one periodically sees it mentioned explicitly (see, e.g., (Hu *et al* 2004) or (Wall 2012, p.8)). It is normally negligible; an important, though theoretical, exception is black hole evaporation, where the presence of gravitons means that black holes evaporate by Hawking radiation even in pure (vacuum) quantum gravity.)

If we go to higher order, we include matter-graviton interactions, and hence indirectly self-gravitational effects of fluctuations in the matter field. It is at this level (again, often disregarding gravitons) that we can derive the linearized accounts of fluctuations which we have seen in inflation, and which also suffice to describe more homely fluctuation phenomena like the Page-Geilker experiment. As those fluctuations grow, the background field approach (whilst remaining formally exact) may become more and more computationally intractable, and less and less useful as a description of the physics. Eventually one might need

to abandon it and apply the more general methods of non-equilibrium quantum field theory; in the context of quantum gravity, these methods are known as *stochastic gravity*²⁷ (see (Hu and Verdaguer 2008) for a review of this relatively new, and growing, subject). But even in doing so, we are simply applying a different and more powerful set of approximation techniques to the same quantum field theory.

To conclude this section: we can define quantum gravity non-perturbatively as a quantum field theory, and extract both semiclassical gravity and linearized quantum gravity as approximations to that theory valid in certain regimes, using (at least at the broad schematic level) essentially the same techniques as we use in non-gravitational QFT. The path-integral formulation of low-energy quantum gravity provides a framework that unifies the various apparently-disconnected fragments of quantum gravity we have discussed so far and puts them under a single, albeit non-fundamental, theory.

9 The scope of low-energy quantum gravity

For a theory as complicated as a quantum field theory (like QED, for instance), it rarely makes sense to say that any given experiment confirms the theory as a whole.²⁸ Rather, what we test in applications of these theories is invariably some particular approximation, regime of application, or limiting case to the theory. Some parts of the evidence for QED consist of evidence for its nonrelativistic electrostatic approximation; others support the existence of the photon; still others provide evidence for coherent interference between those photons. Collectively, the evidence for these various applications of QED provide very strong support for the theory as a whole, while always permitting the logical possibility that some different theory could underpin the selfsame class of approximations and idealizations that we actually test.

As an inevitably-partial classification of these approximation schemes for QED, we can consider asking two questions of the regime we want to analyze:

1. Is it relativistic? In the nonrelativistic regime, radiation can be neglected, long-range interactions can be treated as instantaneous, and particle creation and annihilation does not occur. In the relativistic regime, all these phenomena must be allowed for to varying degrees.
2. How quantum-mechanically does it treat the electromagnetic field? We can distinguish (i) a semiclassical regime, where the field is treated as

²⁷The term ‘stochastic gravity’ is actually used both for the constructive approximation to an underlying QFT that I sketch here, and to more explicitly phenomenological approaches that seek to generalise the semiclassical equation by adding noise terms without requiring any principled microtheory. The two approaches are used concurrently by the same authors and I don’t myself see any real conflict between them (they mirror the way phenomenological physics is developed in other areas). However, see (Mattingly 2014) for discussion of stochastic gravity specifically from the phenomenological perspective.

²⁸See Cartwright (1983) for further development of this point (though she draws more sweeping conclusions about the evidence for theories than I would want to endorse).

sourced by the average charge distribution; (ii) a perturbative regime, where we need to treat the electromagnetic field as a superposition of field values but where the different terms in the superposition may be regarded as small fluctuations around a fixed background (normally the zero field); (iii) a macroscopic-superposition regime,²⁹ where we need to allow the quantum state to evolve into a superposition of macroscopically distinct backgrounds. In case (ii) we can further distinguish (iia) a perturbative/incoherent regime, where interference between terms in the superposition can be neglected and the state behaves like a classical mixture; (iib) a perturbative/coherent regime, where interference cannot be so neglected. (In principle we could make the same distinction for (iii); in practice, decoherence normally means that we cannot expect to observe interference between macroscopically distinct terms in a superposition.)

One way to see the strength of the evidence for QED is to note that there is extensive empirical support for all eight regimes derived from the various combinations of answers to these questions. (I don't intend the following list to be remotely exhaustive.) The semiclassical regime for QED leads to the Hartree-Fock approximation, very extensively applied to relativistic and nonrelativistic atoms and molecules. The perturbative/incoherent regime includes Coulombic scattering theory at nonrelativistic energies, and the standard photon/electron scattering theory of QED at relativistic energies. The perturbative/coherent regime includes multi-atom and multi-ion entanglement in nonrelativistic physics, and the classic photon two-slit experiment and its many generalizations in relativistic physics. And spontaneous breaking of symmetry, and its spontaneous onset as systems are cooled below their critical temperature (in, e.g., superconductors or Bose-Einstein condensates) requires the macroscopic-superposition regime, including (when we want to consider electromagnetic radiation incident on a superconductor, for instance) relativistic applications of that regime. (See Table 1 for a summary.)

Performing the same exercise for low-energy quantum gravity (table 2) shows, unsurprisingly, that it is less well confirmed empirically than QED. (See Table 2 for a summary.) We have really abundant evidence from astrophysics and cosmology for the semiclassical regime both at nonrelativistic (section 2) and relativistic (section 3) energies, pretty good evidence (section 5) for the nonrelativistic part of the perturbative/incoherent regime from the Page-Geilker experiment and more general considerations of Solar-system-scale physics, and some evidence (section 6) for the relativistic part of that regime from primordial cosmology. Stochastic gravity methods allow us to probe the macroscopic regime, mostly in the context of early-universe structure formation (and the boundary between those methods and perturbative treatments of fluctuation

²⁹I don't intend to require any particular stance on the quantum measurement problem here. These macroscopic superpositions, at least for the purposes of this paper, can be interpreted *a la* Everett, but can also be thought of as probability distributions over actually-occurring macroscopic outcomes, where those probability distributions are interpreted via some unspecified hidden-variable theory or dynamical collapse mechanism or as part of a broader Copenhagen-style reading of quantum mechanics.

Table 1: Evidence for QED in various regimes

Quantum superpositions	Non-relativistic	Relativistic
Semiclassical	Hartree-Fock methods (nonrelativistic atoms)	Hartree-Fock methods (relativistic atoms)
Perturbative / Incoherent	non-relativistic (Coulomb) scattering	electron-photon scattering
Perturbative / Coherent	multi-atom and multi-ion entangled states	photon two-slit experiment
Macroscopic	symmetry breaking on cooling: Bose-Einstein condensates	symmetry breaking on cooling: superconductors

formation is not completely sharp) but to my knowledge have not yet produced really robust predictions. (This regime only really makes sense relativistically: in the domain of nonrelativistic gravity, perturbative methods seem likely to capture the full content of the theory.)

At present, we have little evidence supporting applications of LEQG in the perturbative/coherent regime, which is significant because it is this regime which distinguishes unitary quantum gravity from variants in which superpositions of mass distributions cause wavefunction collapse, as has been advocated by Dósi (1987, 2014), Penrose (1969, 1996, 2014), Stamp (2015) and others. For that very reason there has been considerable interest in probing that regime, at least nonrelativistically, by attempting to place a relatively-massive object in a superposition and to test the coherence of that superposition via interference experiments. Those experiments are challenging, but not impossible with current technology: recent experimental work by Donadi *et al* (2021) appears to rule out at least the simplest form of the Diósi-Penrose model of gravitational collapse, and in doing so to corroborate at least a part of the perturbative/coherent nonrelativistic regime. We can optimistically hope for LEQG to be more systematically confirmed — or falsified — in that regime in the comparatively near future, with recent proposals for experiments on self-gravitating entangled systems (Marletto and Vedral 2017; Christodoulou and Rovelli 2018). Probing more strongly relativistic aspects of the perturbative/incoherent regime looks far more challenging (performing the two-slit experiment with gravitons, for instance, is ridiculously beyond any current or foreseeable gravity-wave-detection

Table 2: In which regimes is low-energy quantum gravity observationally supported?

Quantum superpositions	Non-relativistic	Relativistic
Semiclassical	strongly supported (planets, white dwarfs)	strongly supported (stars, neutron stars, supernovae, early-universe cosmology)
Perturbative / Incoherent	strongly supported (Page-Geilker, general considerations of solar-system physics)	some support (quantum-fluctuation models of the CMB)
Perturbative / Coherent	limited support so far but experiments are ongoing	no support and little prospect of direct testing
Macroscopic	not applicable (this regime can be analyzed perturbatively without loss)	developing (stochastic-gravity methods probe this regime but no decisive confirmed predictions)

technology) so that it seems likely any tests of this regime will have to be indirect, probably via cosmology. And no breakdown of this kind can hope to be exhaustive: no doubt the theory will sooner or later be put to the test in regimes that defeat my categorization.

In summary: we know how to apply LEQG in a fairly wide, and growing, range of physical situations. In some of those situations there is extremely good evidence for its accuracy; in others there is some evidence; in others there is the prospect of getting evidence in the near future; in others still, there is little prospect. Certainly it would be premature to treat it as confirmed to anything like the degree of QED; certainly, it is well worth striving to confirm it further observationally or experimentally, in novel regimes. And the very concreteness of the predictions that the theory makes serve to sharpen the task of any such observation or experiment. Low-energy quantum gravity shows us that it is too pessimistic to suppose that quantum gravity must remain a theoretical endeavour: if the predictions of LEQG are often hard to test, still they are made at energy scales far more accessible than the Planck scale. Still, the theory provides a reasonably satisfactory unifying basis for all those quantum-gravitational applications we currently have and in doing so is reasonably well confirmed by experiment. And — perhaps with the exception of the various gravitational-induced collapse models — it is difficult readily to see what alternative could provide a unified underpinning for those parts of LEQG which are strongly supported by evidence.

For all that, it is not fundamental. We have strong reasons to believe that the theory will fail at sufficiently high energies, and that it has other issues — notably the infamous Cosmological Constant Problem — that may point to problems over and above that failure. I discuss the latter in section 11; before that, I discuss the limitations on LEQG that arise from its nature as a nonrenormalizable effective field theory.

10 Quantum gravity at *high* energies

Any quantum field theory is defined with an (implicit or explicit) high-energy cutoff. At the formal level, the cutoff serves to make the path integral (tolerably) well-defined, and to regulate the divergences in perturbative calculations. At the physical level, the cutoff represents a division between those degrees of freedom that the QFT explicitly models and those which it does not. There is in general no requirement that the cutoff represents a physical scale at which the physics fundamentally changes: given a QFT with a cutoff at some energy E , we can always integrate out degrees of freedom between λE and E and obtain a new QFT with a cutoff at energy λE . But if we ask the inverse question — how *high* can the cutoff be set before the QFT breaks down? — then we get different answers for different QFTs.

The best behaved cases are *asymptotically safe* QFTs, of which the physically most significant examples are pure non-Abelian gauge theories like QCD. Here the cutoff can be set arbitrarily high, and the interaction strengths in the theory

become lower as particle energies increase, tending to zero or to some finite value. For these QFTs it is at least a live possibility to suppose that there is a well-defined continuum limit of the theory which could be defined in a mathematically rigorous way (though the physical significance of doing so can be questioned; cf (Wallace 2011)).

Next there are renormalizable but non-asymptotically-free theories; here some interactions actually increase in strength at higher energies, and there is good reason to suspect that *very* high values of the cutoff eventually cause the interaction strengths to diverge and the theory to become ill-defined. However, the energy scales on which these ‘Landau poles’ occur are higher by far than any physically relevant energy scales (higher by far than the Planck scale, for instance). QED and (the $U(1)$ and Higgs sectors of) the Standard Model are of this form: we have reason to doubt that they have a genuine continuum limit, but they can be well defined at energy levels far beyond those for which external reasons tell us to expect them to lose empirical validity.

The least well behaved cases are *nonrenormalizable* theories. Here the theory imposes a much lower and sharper bound on the value of the cutoff. In a nonrenormalizable theory, if we try to do calculations that approach that bound, we will first lose predictive power (each new order of perturbative calculation will involve new constants, knowable only from experiment, whose effects are not sharply suppressed unless we calculate at energies far below the bound). If we persist, unitarity itself will break down. For theories of this kind, the bound on the cutoff represents a directly physical, internal-to-the-theory constraint on the energies at which it is applicable. Various low-energy sectors of the Standard Model — notably weak-interaction-induced lepton scattering and pion theory, as well as the electroweak sector as a whole without the Higgs boson — are well described by nonrenormalizable QFTs; more importantly for our purposes, low-energy quantum gravity is non-renormalizable.

In the days before the effective-field-theory paradigm, the cutoff was widely regarded as a purely formal device, ideally to be set to infinity at the end of a calculation. From this perspective,³⁰ the only really satisfactory QFTs are the asymptotically-free ones. Renormalizable theories with Landau poles can be tolerated but induce unease. And non-renormalizable theories are pathological and must be disregarded.

But if a quantum theory is thought of as an effective field theory, nonrenormalizable theories lose their stigma; instead, they convey valuable information about the energy levels at which new physics can be expected. The four-fermion theory of weak lepton scattering must be cut off at energies of $\sim 100\text{GeV}$, pointing to new physics at around that scale; that ‘new physics’ turns out to be the W and Z bosons which mediate the weak interaction, with masses of 80GeV and 92GeV respectively. (And the creation of these particles in high-energy scattering of leptons provides a directly physical explanation of unitarity viola-

³⁰It is probably fair to say that ‘this perspective’ is a little over-simplified as a description of the actual history: physicists in the pre-effective-field-theory era had more subtle and nuanced views of the high-energy cutoff than modern physics textbooks might suggest. See (Li 2013) for further discussion.

tion in the lepton sector). The Higgs-free part of the Standard Model must be cut off at energies of a few hundred GeV to be well-behaved; hence, whether or not the Higgs boson was found in the LHC, physicists could be confident that *some* new physics would be found.³¹ And LEQG must be cut off at the Planck length (and is computationally reliable only for energies well below the Planck length), which tells us that new physics will be needed to describe systems at Planckian energy scales.³²

There is also a more directly physical argument why LEQG, and the methods of QFT in general, ought to fail at around the Planck length. (Here I follow Susskind and Lindesay (2005, pp.95-6).) In QFT, the effective minimum size $r(m)$ of an quantum excitation of mass m is set by its Compton wavelength — $r(m) \geq 1/m$, in the units we are using. But in general relativity, there is another constraint on the minimum size of a particle: it must lie outside its own Schwarzschild radius, or else an event horizon will form around it and (at least as regards external physics) it will behave as a black hole of that radius. This constraint imposes $r(m) \geq 2Gm \sim (l_P/m_P)m$, where l_P and m_P are the Planck length and mass, respectively (note that as we should expect from this classical calculation, their ratio does not depend on \hbar ,).

These constraints cross over at the Planck scale: a particle with a mass greater than $\sim m_P$ has a Compton wavelength contained within its own Schwarzschild radius. The Planck length is thus a cutoff on the energies of elementary (i. e., non-compound) excitations of quantum fields. If a spacetime description is applicable at all at scales below the Planck length, it cannot be a description in terms of excitations of quantum fields, of the kind we are used to from QFT.

We can see this in another way by considering scattering theory. In QFT, the energies of the products of a collision with center-of-mass energy m also have energy $\sim m$. But if we collide elementary particles with a center-of-mass energy $> m_P$, they will form a black hole, which will decay by emitting Hawking quanta of energy $\sim m_P/m$. In other words, above the Planck length the energy of collision products actually begins to decrease again. As Susskind and Lindesay put it (pp.95-6):

Thus we see that a giant “Super Plancktron Collider” (SPC) would fail to discover fundamental length scales smaller than the Planck scale l_P , no matter how high the energy. In fact, as the energy increased we would be probing still larger scales. . .

What can we say about quantum gravity at Planckian scales, given this failure of LEQG at those scales? Let’s distinguish two sorts of theory:

³¹For a review of the physics here, see (Weinberg 1995, chs.19,21) or (Donoghue *et al* 2014).

³²This is the standard view in physics on renormalizability of LEQG; there is, however, a longstanding research program (‘asymptotic safety’) to explore whether non-perturbative methods might after all make the theory definable at all scales. (The original idea is due to Weinberg (1979); see (Niedermaier 2007; Reuter and Saueressig 2012) for reviews). For my purposes this caveat doesn’t matter much: as I explain below, even if there is an extension of LEQG in this way, its interpretation at Planckian scales will deviate wildly from the relatively-tame QFT interpretation of LEQG.

1. An *ultraviolet quantum gravity* theory (UVQG) is any theory from which LEQG can be derived as an effective field theory valid in some limited regime of the UVQG.
2. A *correct quantum gravity* theory is any theory which actually gets the physics right in some regime which includes (i) the domain of empirical validation of LEQG and (ii) at least some part of the Planckian regime.

The two are not logically equivalent. Empirical adequacy for a quantum theory of gravity, at present, only requires it to reproduce the predictions of LEQG in certain domains: the semiclassical domain, the perturbative/incoherent domain; soon, perhaps, the nonrelativistic part of the perturbative/coherent domain. (One sometimes hears it said that the theory need only reproduce the predictions of non-gravitational QFT and classical GR; sections 2–6 should make clear why this is far too weak a demand to make.) Most physicists working on quantum gravity set themselves the more demanding task of reproducing most or all of LEQG, notably including its predictions of black hole evaporation, which are unsupported by any observational evidence to date: that is, they are searching for a UVQG theory. In doing so, they go beyond what evidence requires of them. That said, their strategy seems well-motivated pending a low-energy alternative to LEQG. The problem of reproducing those fragments of LEQG that are directly supported by evidence is just too open to provide much guidance, whereas reproducing all of LEQG is a very strong constraint on theorizing.

Conversely, there is no obvious reason why a UVQG need be a (let alone the) *correct* quantum gravity theory. If we had two inequivalent ultraviolet completions of LEQG, nothing in LEQG would allow us to distinguish them. And since we have virtually no observational access to regimes in which LEQG is not valid, prospects would be dim for any means of distinguishing the two. However concerning this might be, however, it is a rather theoretical problem at the current state of development of quantum gravity, where our problem is not that we have two or more UVQGs but that we do not have even one that is generally accepted to be successful in reproducing LEQG in all the domains in which it can be applied.

As a consequence, and notwithstanding these constraints, when physicists speak of a ‘quantum theory of gravity’ they normally mean an ultraviolet completion of LEQG (when they do not mean LEQG itself). And the basic requirement that the theory *is* an ultraviolet completion of LEQG does tell us at least a little about it.

Firstly, it is not a QFT, or at any rate not a QFT interpretable in anything like the way we normally interpret QFTs (including LEQG). The ‘excitations on a background spacetime’ picture of QFTs, as we have seen, breaks down at Planckian scales; well below those scales, LEQG works fine and there is no obvious extra work for a UVQG to do. If QFT methods remain applicable for Planck-scale physics, they must be applicable in a radically different way. (The AdS/CFT correspondence, arguably our best candidate so far for UVQG, illustrates how this might happen: quantum-gravity phenomena in the interior

of asymptotically-AdS spacetime are describable indirectly via a QFT on the conformal boundary of that spacetime.)

Secondly, by the same token a UVQG will be ultraviolet *complete*, and not just another effective field theory. The reason is not that we hubristically expect our theory to apply at arbitrarily high energy scales, but that the idea of an energy scale in the normal way we use it loses meaning beyond the Planck scale. To say that a QFT is ultraviolet-complete, mathematically, is to say (among other things) that the Fourier-transformed four-point function remains well-behaved at arbitrarily large values of input wavenumber. But the normal physical interpretation of that four-point function as encoding particle scattering cross-sections at successively higher energies becomes meaningless at Planckian scales. If our UVQG is restricted in its domain of application, that restriction will need to be stated in novel ways and not by our usual considerations of energy scale.³³ (Again, AdS/CFT correspondence offers an illustration of how this could occur: the ‘asymptotically AdS’ boundary conditions are normally interpreted as a way of putting a system in a box, so that AdS/CFT gives a quantum-gravity description of isolated systems but is inadequate for cosmology.)

Thirdly, it will need to be a ‘theory of everything’ in the sense that it will need to incorporate the non-gravitational degrees of freedom described by the Standard Model, not just purely-gravitational degrees of freedom. The reason is comparatively mundane: there is just no clean separation of these degrees of freedom at Planckian scales. The annihilation of two gravitons to produce leptons and antileptons, for instance, is a risibly unimportant process in LEQG, but becomes significant as we approach the Planck scale. There is no reason to expect a quantum theory that reproduces just the vacuum sector of LEQG to be empirically adequate. This is, however, neutral as to whether a UVQG includes a more thoroughgoing unification of fundamental forces (as string theory attempts to do) or more conservatively maintains the gauge structure of the Standard Model (as loop quantum gravity does).

Finally, and most importantly, it will be highly underconstrained by observational evidence except insofar as that evidence is evidence for LEQG. This is what is meant when experts talk about the near-absence of evidence for quantum gravity: it is not that they are ignorant of LEQG but that they mean ultraviolet quantum gravity, and are speaking of evidence that goes beyond evidence for LEQG. And that is, indeed, difficult in the extreme to obtain, requiring energy levels which seem to occur in our Universe only behind the screen of the plasma of the very-early Universe or the still more opaque screen of a black hole’s event horizon. From this point of view, the success of LEQG stands to the search for UVQG much as the success of the Standard model stands to the search for post-Standard-Model physics: we have excellent grounds to think that the theory is not the last word, but nonetheless it is so successful as to make it difficult to go beyond it.

So it is exciting as well as puzzling that LEQG appears to contain a severe

³³Though see (Crowther and Linnemann 2017) for a dissenting view.

anomaly that might point towards new physics: the mysteriously low value of the observed cosmological constant.

11 The cosmological constant problem

Suppose that we want to apply low energy quantum gravity in the background-field regime and for some comparatively tame region of the background spacetime. Specifically, let's suppose that the scale of spacetime curvature in this region is $\gg 10^{-15}\text{m}$, that the matter field can be described in this region as a small perturbation away from some thermal-equilibrium state, and that that state is either the vacuum or at any rate is at a temperature $\ll 10^{14}\text{K}$. (These are not unduly demanding constraints. The curvature scale near the event horizon of an astrophysical black hole is $\sim 10^4\text{m}$; the temperature in the core of a supergiant star just before becoming a supernova is $\sim 10^{12}\text{K}$.) A particle of energy $> 100\text{MeV}$ in this region will effectively behave as if in the Minkowski vacuum (the Compton wavelength of a 100-MeV particle is $\sim 10^{-15}\text{m}$; the typical energy of a particle in an equilibrium state at temperature 10^{14}K is $\sim 100\text{MeV}$). Or put another way, at least for excitations on scales of $\sim 100\text{MeV}$ or higher, we should be able to make local calculations as if in the Minkowski vacuum.³⁴

Let's now try to apply the rules of the background-field method: for self-consistency (and ignoring higher-order terms in the gravitational action), we need

$$\frac{1}{8\pi G}(G_{\mu\nu} + g_{\mu\nu}\Lambda_0) = \langle \widehat{T}_{\mu\nu} \rangle \quad (30)$$

where the expectation value is calculated for the quantum state of the matter fields (and, in principle, the gravitons) on the background. What we would hope to find is that there are two contributions to that stress-energy expectation value: one due to the thermal energy of the finite-temperature local equilibrium state (if we are indeed at finite temperature) and one due to quantum excitations. The latter term should be treatable as a sum of particle energies if the excitations are on lengthscales short compared to the curvature.

We do indeed find that, but we also find something much less welcome. For the moment, let's suppose that the calculation is carried out on a Minkowski background. We then get a series of very large terms proportional to the Minkowski metric. By far the largest of these terms, is of order $(E_{\text{cutoff}})^4$, where E_{cutoff} is the cutoff energy scale we are using to regularize LEQG (normally assumed to be around the Planck scale, $\sim 10^{19}\text{GeV}$); this term can be understood as the zero-point energy density of the vacuum. The second term is of order $(E_{EW})^4$, where $E_{EW} \sim 100\text{GeV}$ is the scale of electroweak symmetry breaking; it can be understood as the energy density of the Higgs condensate,

³⁴Why not the de Sitter vacuum, given the positive observed cosmological constant? You can use that if you like, but it doesn't matter (and is mildly more annoying technically). The observed value of the cosmological constant corresponds to a curvature scale of $\sim 10^{26}\text{m}$, or $10^{38} \times$ the Compton wavelength of our 100-MeV particle. The local physics just doesn't care about the difference between this level of curvature and exact flatness, not unless we're looking at cosmological-wavelength excitations.

due to the spontaneous breaking of electroweak symmetry. The third term is of order $(E_{CH})^4$ where $E_{CH} \sim 1\text{GeV}$ is the scale of approximate chiral symmetry breaking in the QCD sector of the Standard Model. (Our stricture that $T \ll 10^{14}\text{K}$ ensures that these broken symmetries remain broken for our quantum state.)

(As a calculational matter, the two condensate terms occur at tree order in the calculation, and the zero-point term occurs at one-loop order. There will be one-loop quantum corrections to the condensate terms, and higher-order corrections to all the terms, but they do not change the basic magnitudes of the terms. They do somewhat blur the clean distinction between the three contributions, but not in a way that really matters for the argument here.)

The zero-point term arises due to fluctuations on scales far smaller than the assumed curvature of our local region, and the other two terms are due to potential effects insensitive to curvature, so we can take our Minkowskian calculation over to that region just by replacing the Minkowski metric with the local background metric. The semiclassical equation can then be written as

$$\frac{1}{8\pi G}(G_{\mu\nu} + g_{\mu\nu}\Lambda_0) = \left\langle \widehat{T}_{\mu\nu}^{\text{IR}} \right\rangle + \frac{1}{8\pi G}g_{\mu\nu}(\Lambda_{ZP} + \Lambda_{EW} + \Lambda_{CH}) \quad (31)$$

or (rearranging)

$$\frac{1}{8\pi G}(G_{\mu\nu} + g_{\mu\nu}(\Lambda_0 - \Lambda_{ZP} - \Lambda_{EW} - \Lambda_{CH})) = \left\langle \widehat{T}_{\mu\nu}^{\text{IR}} \right\rangle \quad (32)$$

where $\Lambda_{ZP}/8\pi G$, $\Lambda_{EW}/8\pi G$, and $\Lambda_{CH}/8\pi G$ are respectively the zero-point, Higgs condensate, and chiral condensate contributions to the stress-energy density, and where (translating into SI units for clarity),

$$\begin{aligned} \Lambda_{ZP} &\sim 10^{97}\text{kg m}^{-3} \\ \Lambda_{EW} &\sim 10^{29}\text{kg m}^{-3} \\ \Lambda_{CH} &\sim 10^{21}\text{kg m}^{-3}. \end{aligned} \quad (33)$$

The first term on the right-hand-side of equation (31), the ‘infrared’ part of the stress-energy tensor, encodes contributions from the finite-temperature background, from particle excitations on scales small compared to the curvature, and potentially from other excitations on longer scales to which the particle concept cannot be applied; however, on basic dimensional-analysis grounds it will be far smaller than any of the other three terms.

Now, the observationally-supported version of the semiclassical equation is

$$\frac{1}{8\pi G}(G_{\mu\nu} + g_{\mu\nu}\Lambda_{obs}) = \left\langle \widehat{T}_{\mu\nu}^{\text{IR}} \right\rangle \quad (34)$$

where Λ_{obs} is the observed cosmological constant. Comparing equations (32) and (34), we get

$$\Lambda_{obs} = \Lambda_0 - (\Lambda_{ZP} + \Lambda_{EW} + \Lambda_{CH}). \quad (35)$$

In one sense, there is no difficulty solving this equation. Λ_0 — the *bare* cosmological constant, in particle-physics terminology — influences observable quantities only through equation (35), and so we can take (35) simply as fixing the bare cosmological constant as a function of the observable parameters and the high-energy cutoff, as usual in QFT.

On the other hand,

$$\Lambda_{obs} \sim 10^{-26} \text{kg m}^{-3}, \quad (36)$$

which is about 120 orders of magnitude smaller than the zero-point energy density and about 50 orders of magnitude smaller even than the chiral-condensate energy density. So the small value of the observed cosmological constant we observe today requires the bare cosmological constant — or, equivalently, the bare parameters of the theory as a whole — to be fine-tuned to an extraordinarily precise value, to one part in 10^{120} or more.³⁵ This is the cosmological constant problem (CCP).³⁶

Whether fine-tuning to this degree is a problem at all, and if so how it might be resolved, lies beyond the scope of this article. Maybe we should be relaxed about it, since after all Λ_0 is unobservable; maybe we can solve it with subtle short-distance/long-distance effects which correct the effective-field-theory framework; maybe a resolution needs to draw on anthropic arguments or something still more radical.³⁷ But the fine-tuning arises as a direct result of comparatively simple calculations within LEQG.

Recent critical discussions of the CCP (e.g., Saunders 2002; Koberinski 2017, 2021; Schneider 2020) have generally presented the problem quite differently, as relying on multiple conceptual assumptions and speculations about as-yet-undeveloped physics. On this approach, we first have to establish that the zero-point energy is real at all, perhaps by consideration of the Casimir effect; then we have to argue that it gravitates; then we have to argue that the way it gravitates is via the semiclassical equation. At each step, it is suggested, the problem turns on some specific conceptual choices, and by the end of the story one is left with the impression that the existence of the CCP is speculative at best. But if the CCP is thought of as a problem inside LEQG, no conceptual steps are required. All one has to do is derive the semiclassical equation via the background-field method and then calculate the terms in it.

(That calculation can itself be challenged. Koberinski (2021) argues that the calculated value of the zero-point energy, depending sensitively on the cutoff as it does, cannot be treated as a physical prediction of the theory; Koberinski and Smeenk (in forthcoming work) also raise concerns about the applicability

³⁵The relationships between bare and observed parameters in particle physics are *not* generally fine-tuned in this way; the exception is the Higgs mass, where fine-tuning concerns constitute the closely-related *hierarchy problem* of the Standard model.

³⁶The literature on the CCP is too vast to systematically reference here. The classic statement of the problem remains (Weinberg 1989); more recent introductions/reviews include (Carroll 2001; Polchinski 2006; Burgess 2013; Padilla 2015); (Rugh and Zinkernagel 2002) is a good non-technical introduction.

³⁷For scepticism about the significance of the fine-tuning see (Bianchi and Rovelli 2010; Hossenfelder 2018); for anthropic considerations see (Polchinski 2006); see also (Wallace 2019b) and references therein.

of the ‘Minkowskian approximation’ which I appealed to above when applied to even gently curved spacetimes. I find both objections unpersuasive. As to the first: the fact that the zero-point energy is extremely large is *not* sensitive to the cutoff; nor is the fact that the bare cosmological constant requires extreme fine-tuning. And we cannot simply drop a term in a calculation because we find it unreliable, not least because we need to extract a finite, physically-significant temperature-dependent component from the zero-point term. As to the second: while it is certainly worth being concerned about various ways in which QFT calculations might subtly fail, the Minkowskian method is used pretty widely in calculating equations of state in astrophysics and cosmology, and more generally, the basic idea of the separation-of-scales approach which lets us treat short-wavelength phenomena as insensitive to long-distance features of the background is totally ubiquitous in physics. But in any case, we can sidestep these concerns entirely just by appealing to the Higgs-condensate and chiral-condensate contributions to the stress-energy density, which are insensitive to the high-energy cutoff and can be calculated at tree level without needing to appeal to any Minkowskian approximation to the field modes.)

If the CCP is a consequence of low energy quantum gravity, doesn’t it give us reason to worry that something is wrong with that theory? And couldn’t we avoid the CCP just by abandoning LEQG? Yes, and yes. The latter is trivial: *of course* we can solve a deep conceptual problem in our best current theory, whatever it is, by discarding it. (Throw away quantum theory and *voilà*, no quantum measurement problem!) The former is pretty much universally recognized in the physics community: Weinberg’s famous paper on the cosmological constant problem (Weinberg 1989) begins “physics thrives on crisis” and analogizes the CCP to other historical crises that have led to new physics. But a successful resolution of the CCP requires not just the bare observation that a theory other than LEQG might not have a CCP, but an actual construction of such a theory so as to preserve the observational successes of LEQG. And the challenge is all the more stark because the problem arises already in the semiclassical regime of LEQG, which is exactly the regime where observational evidence is strongest.

We can state the CCP as a simple challenge. Suppose that you are tasked with studying astrophysics or cosmology in a regime where a quantum-field-theoretic description of matter is needed (but where fluctuations in the matter distribution are not significant): perhaps the regime is the interior of neutron stars; perhaps the core of a supernova; perhaps the quark/gluon plasma of the reasonably-early Universe. To make any progress you will need to calculate an equation of state for the matter. If you calculate the equation of state using the methods of quantum field theory to derive the expected stress-energy tensor, you will immediately encounter very large contributions proportional to the metric from the Higgs and chiral condensates, and will encounter even larger contributions once you proceed beyond leading order, and you will need to cancel them out using a very precisely tuned bare cosmological constant in order to match observed results; in other words, you will be confronted with the CCP. If you plan to perform your calculation without using QFT in this way,

then you had better be able to say what your plan is.

Kuhn (1962) famously argued that scientific theories are replaced not because clear experiments falsify them, but through the slow build-up of ‘anomalies’ in our extant theories — theoretical or empirical problems in the theory. Confronted with an anomaly, practicing scientists do not simply abandon an otherwise-successful theory; they note the anomaly and work around it. Sometimes anomalies are resolved through deeper understanding of the theory, but at other times, they point the way to an eventual replacement of the theory. The CCP is a Kuhnian anomaly in low-energy quantum gravity: absent some better theory, it is no reason to abandon LEQG, but it points to deficiencies either in LEQG or in our understanding of it, and may help us obtain that better theory.

12 Conclusion

If ‘quantum gravity’ means ‘a physical theory which combines general relativity and quantum theory’, or ‘a physical theory which has classical general relativity and non-gravitational QFT as limiting cases or partial approximations’ then we have such a theory, and have had it for decades. It is the straightforward, conservative theory one obtains by treating GR as an effective field theory, quantized through path-integral means. There is much we do not understand about even this theory, but we understand enough to know that it is applicable to a wide variety of observational contexts, many of which are accessible with current observational and experimental methods. Indeed, the theory has been widely applied in many such contexts, and those applications match the observational data; those applications are invariably through one or more approximation schemes, but that is true of virtually any application of a complicated physical theory to concrete examples. Much (most?) of contemporary astrophysics and cosmology relies on our ability to synthesise quantum-mechanical and gravitational phenomena, and low-energy quantum gravity provides such a synthesis which seems fairly satisfactory in its own terms; its nonrenormalizability, and its inevitable failure at Planckian scales, is unexceptional in the era of effective field theories. And since it includes as limiting cases both (semi)classical and phenomenological general relativity, and the Standard Model of particle physics, it has a reasonable claim to be the most widely applicable physics theory ever obtained, with a scope covering almost every experiment and observation we could hope to make. But by the same token, the theory makes many predictions which have yet to be tested but which lie well below the high-inaccessible Planck scale.

If ‘quantum gravity’ means instead ‘an ultraviolet completion of low-energy quantum gravity’, or ‘a theory of gravity that applies at Planckian scales’, then it is true that we do not have such a theory, and also true that very little empirical evidence bears on such a theory — except insofar as it also bears on low-energy quantum gravity. As such, reproducing low-energy quantum gravity — and not simply non-gravitational QFT and classical general relativity — is generally taken as the success condition for any such theory. That does not, and should

not, rule out attempts to construct a Planck-scale quantum gravity theory which differs from low-energy quantum gravity in a yet-unexplored regime. But any such attempts still have large fragments of low-energy quantum gravity, going well beyond non-gravitational QFT and classical general relativity — that they need to reproduce to match observations.

The Cosmological Constant Problem arises within low-energy quantum gravity via straightforward calculation: the parameters in the theory need to be tuned to a precision of one part in 10^{120} to reproduce the observed cosmological constant. This does not falsify low-energy quantum gravity, but it is hard to explain that fine-tuning within the established paradigm of effective field theory. The cosmological constant problem is thus one of the few clues we have as to what might lie beyond this extraordinarily powerful and successful, yet ultimately limited, theory.

Acknowledgements

Thanks to Adam Koberinski, Simon Saunders, Mike Schneider, and audiences in Oxford and Pittsburgh for useful discussion; thanks also to Adam for comments on a previous draft.

References

- Ade, P. et al. (2018). Constraints on primordial gravitational waves using Planck, WMAP and new BICEP2/Keck observations through the 2015 season. *Physical Review Letters* 121, 221301.
- Arnett, W. D. (1996). *Supernovae and nucleosynthesis: an investigation of the history of matter, from the Big Bang to the present*. Princeton: Princeton University Press.
- Ashby, N. (2003). Relativity in the Global Positioning System. *Living Reviews in Relativity* 6, 1.
- Banks, T. (2008). *Modern Quantum Field Theory: a Concise Introduction*. Cambridge: Cambridge University Press.
- Belot, G. (2018). Fifty million Elvis fans can't be wrong. *Nous* 52, 946–981.
- Berges, J. (2004). Introduction to nonequilibrium quantum field theory. *AIP Conference Proceedings* 739, 3.
- Bhattacharya, K., S. Mohanty, and A. Nautiyal (2006). Enhanced polarization of the cosmic microwave background radiation from thermal gravitational waves. *Physical Review Letters* 97, 251301.
- Bianchi, E. and C. Rovelli (2010). Why all these prejudices against a constant? <https://arxiv.org/abs/1002.3966>.
- Binney, J. and S. Tremaine (2008). *Galactic Dynamics* (2nd ed.). Princeton: Princeton University Press.

- Bonitz, M. (2016). *Quantum Kinetic Theory* (2nd ed.). Heidelberg: Springer.
- Burgess, C. (2004). Quantum gravity in everyday life: General relativity as an effective field theory. *Living Reviews in Relativity* 5, 7.
- Burgess, C. (2013). The cosmological constant problem: Why it's hard to get dark energy from microphysics. <https://arxiv.org/abs/1309.4133>.
- Butterfield, J. and C. Isham (2001). Spacetime and the philosophical challenge of quantum gravity. In C. Callender and N. Huggett (Eds.), *Physics Meets Philosophy at the Planck Scale: Contemporary Theories in Quantum Gravity*, pp. 33–89. Cambridge, UK: Cambridge University Press.
- Butterfield, J. and C. J. Isham (1999). On the emergence of time in quantum gravity. In J. Butterfield (Ed.), *Arguments of Time*, pp. 111–168. Oxford University Press.
- Callender, C. and N. Huggett (2001). Introduction. In C. Callender and N. Huggett (Eds.), *Physics meets philosophy at the Planck scale: Contemporary theories in quantum gravity*, pp. 1–32. Cambridge, UK: Cambridge University Press.
- Calzetta, E. A. and B.-L. B. Hu (2008). *Nonequilibrium Quantum Field Theory*. Cambridge: Cambridge University Press.
- Carroll, S. (2001). The cosmological constant. *Living Reviews in Relativity* 4, 1.
- Carroll, S. (2003). *Spacetime and Geometry: an Introduction to General Relativity*. San Francisco, CA.: Addison Wesley.
- Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- Christodoulou, M. and C. Rovelli (2018). On the possibility of laboratory evidence for quantum superposition of geometries. *Physics Letters B* 792, 94–98.
- Coleman, S. and E. Weinberg (1973). Radiative corrections as the origin of spontaneous symmetry breaking. *Physical Review D* 7, 1888.
- Coleman Miller, M. and N. Yunes (2019). The new frontier of gravitational waves. *Nature* 568, 469–476.
- Crowther, K. and N. Linnemann (2017). Renormalizability, fundamentality, and a final theory: The role of UV-completion in the search for quantum gravity. *British Journal for the Philosophy of Science* 70, 377–406.
- Diósi, L. (1987). A universal master equation for the gravitational violations of quantum mechanics. *Physics Letters* 120, 377–381.
- Diósi, L. (2014). Gravity related spontaneous wavefunction collapse in bulk matter. *New Journal of Physics* 16, 105006.
- Donadi, S., K. Piscicchia, C. Curceanu, L. Diósi, M. Laubenstein, and A. Bassi (2021). Underground test of gravity-related wavefunction collapse. *Nature Physics* 17, 74–78.

- Donoghue, J. F. (1997). Introduction to the effective field theory description of gravity. In F. Cornet and M. J. Herrero (Eds.), *Proceedings of the Advanced School on Effective Field Theories, Almunecar, Spain*. World Scientific. arxiv.org/gr-qc/9512024.
- Donoghue, J. F. (2012). The effective field theory treatment of quantum gravity. *AIP Conference Proceedings* 1483, 73.
- Donoghue, J. F., E. Golowich, and B. R. Holstein (2014). *Dynamics of the Standard Model (Second Edition)*. Cambridge: Cambridge University Press.
- Fuchs, C. and A. Peres (2000). Quantum theory needs no “interpretation”. *Physics Today* 53(3), 70–71.
- Fuchs, C. A., N. D. Mermin, and R. Schack (2014). An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics* 82, 749–754.
- Gell-Mann, M. and J. B. Hartle (1989). Quantum Mechanics in the Light of Quantum Cosmology. In W. H. Zurek (Ed.), *Complexity, Entropy and the Physics of Information*, pp. 425–459. Redwood City, California: Westview Press.
- Gell-Mann, M. and J. B. Hartle (1993). Classical equations for quantum systems. *Physical Review D* 47, 3345–3382.
- Griffiths, R. (1993). Consistent interpretation of quantum mechanics using quantum trajectories. *Physical Review Letters* 70, 2201–2204.
- Griffiths, R. B. (1984). Consistent histories and the interpretation of quantum mechanics. *Journal of Statistical Physics* 36, 219–272.
- Halliwell, J. J. (1995). A review of the decoherent histories approach to quantum mechanics. *Annals of the New York Academy of Sciences* 755, 726–740.
- Hawking, S. (1979). The path-integral approach to quantum gravity. In S. Hawking and W. Israel (Eds.), *General Relativity: an Einstein centenary survey*, pp. 746–789. Cambridge: Cambridge University Press.
- Healey, R. (2012). Quantum theory: A pragmatist approach. *British Journal for the Philosophy of Science* 63, 729–771.
- Healey, R. (2017). *The Quantum Revolution in Philosophy*. Oxford: Oxford University Press.
- Hossenfelder, S. (2018). Screams for explanation: Finetuning and naturalness in the foundations of physics. <https://arxiv.org/abs/1801.02176>.
- Hu, B., A. Roura, and E. Verdaguer (2004). Induced quantum metric fluctuations and the validity of semiclassical gravity. *Physical Review D* 70, 044002.
- Hu, B. and E. Verdaguer (2008). Stochastic gravity: Theory and applications. *Living Reviews in Relativity* 11, 3.

- Keller, J., C. Wellenhofer, K. Hebeler, and A. Schwenk (2021). Neutron matter at finite temperature based on chiral effective field theory interactions. *Physical Review C* 103, 055806.
- Kiefer, C. (2007). *Quantum Gravity* (2nd ed.). Oxford, UK: Oxford University Press.
- Kippenhahn, R., A. Weigert, and A. Weiss (2012). *Stellar Structure and Evolution* (2nd ed.). New York/Dordrecht/London: Springer.
- Koberinski, A. (2017). Problems with the cosmological constant problem. <http://philsci-archive.pitt.edu/14244/>.
- Koberinski, A. (2021). Regularizing (away) vacuum energy. *Foundations of Physics* 51, 20.
- Kuhn, T. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Li, B. (2013). Interpretive strategies for deductively insecure theories: The case of early quantum electrodynamics. *Studies in History and Philosophy of Modern Physics* 44, 395–403.
- Marletto, C. and V. Vedral (2017). Witness gravity’s quantum side in the lab. *Nature* 547, 156–158.
- Mattingly, J. (2005). Is quantum gravity necessary. In A. Kox and J. Eisenstaedt (Eds.), *The Universe of General Relativity*, pp. 327–328. Boston: Birkhäuser.
- Mattingly, J. (2009). Mongrel gravity. *Erkenntnis* 70, 379–395.
- Mattingly, J. (2014). Unprincipled microgravity. *Studies in History and Philosophy of Modern Physics* 46, 179–185.
- Niedermaier, M. (2007). The asymptotic safety scenario in quantum gravity - an introduction. *Classical and Quantum Gravity* 24, R171–230.
- Omnes, R. (1988). Logical reformulation of quantum mechanics. i. foundations. *Journal of Statistical Physics* 53, 893–932.
- Omnes, R. (1992). Consistent interpretations of quantum mechanics. *Reviews of Modern Physics* 64, 339–382.
- Padilla, A. (2015). Lectures on the cosmological constant problem. <https://arxiv.org/abs/1502.05296>.
- Padmanabhan, T. (2014). A short note on the boundary term for the Hilbert action. *Modern Physics Letters A* 29, 1450037.
- Page, D. N. and C. D. Geilker (1981). Indirect evidence for quantum gravity. *Physical Review Letters* 47, 979–982.
- Penrose, R. (1969). Gravitational collapse: The role of general relativity. *II Nuovo Cimento Numero Speciale I*, 257. Reprinted in *General Relativity and Gravitation* 34 (2002) pp.1141–1165.

- Penrose, R. (1996). On gravity's role in quantum state reduction. *General Relativity and Gravitation* 28, 581–600.
- Penrose, R. (2014). On the gravitization of quantum mechanics 1: Quantum state reduction. *Foundations of Physics* 44, 557–575.
- Peskin, M. E. and D. V. Schroeder (1995). *An introduction to Quantum Field Theory*. Reading, Massachusetts: Addison-Wesley.
- Pinto-Neto, N. and W. Struyve (2018). Bohmian quantum gravity and cosmology. <https://arxiv.org/abs/1801.03353>.
- Polchinski, J. (2006). The cosmological constant and the string landscape. <https://arxiv.org/abs/hep-th/0603249>.
- Reuter, M. and F. Saueressig (2012). Quantum Einstein gravity. <https://arxiv.org/abs/1202.2274>.
- Rickles, D. (2008). Quantum gravity: A primer for philosophers. In D. Rickles (Ed.), *The Ashgate Companion to Contemporary Philosophy of Physics*, pp. 262–365. New York: Ashgate.
- Rovelli, C. (1996). Relational quantum mechanics. *International Journal of Theoretical Physics* 35, 1637. Substantially revised version at arxiv.org/abs/quant-ph/9609002.
- Rovelli, C. (2008). Quantum gravity. *Scholarpedia* 3, 7117. www.scholarpedia.org/article/Quantum_gravity.
- Rugh, S. E. and H. Zinkernagel (2002). The quantum vacuum and the cosmological constant problem. *Studies in History and Philosophy of Modern Physics* 33, 663–705.
- Saunders, S. (2002). Is the zero-point energy real? In M. Kuhlmann, H. Lyre, and A. Wayne (Eds.), *Ontological Aspects of Quantum Field Theory*, pp. 313–344. Singapore: World Scientific.
- Schneider, M. D. (2020). What's the problem with the cosmological constant? *Philosophy of Science* 87, 1–20.
- Smeenk, C. (2017). Testing inflation. In K. Chamchan, J. Barrow, and S. Saunders (Eds.), *Philosophy of Cosmology*. Cambridge: Cambridge University Press.
- Stamp, P. (2015). Rationale for a correlated worldline theory of quantum gravity. *New Journal of Physics* 17, 065017.
- Susskind, L. and J. Lindesay (2005). *An introduction to black holes, information and the string theory revolution: the Holographic Universe*. Singapore: World Scientific.
- Urena-Lopez, L. A. (2019). Brief review on scalar field dark matter models. *Frontiers of Astronomy and Space Science*. <https://doi.org/10.3389/fspas.2019.00047>.
- Wald, R. M. (1984). *General Relativity*. Chicago: University of Chicago Press.

- Wall, A. C. (2012). A proof of the generalized second law for rapidly changing fields and arbitrary horizon slices. *Physical Review D* 85, 104049. Most recent version at <https://arxiv.org/abs/1105.3445v5>.
- Wallace, D. (2011). Taking particle physics seriously: a critique of the algebraic approach to quantum field theory. *Studies in the History and Philosophy of Modern Physics* 42, 116–125.
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.
- Wallace, D. (2016a). Interpreting the quantum mechanics of cosmology. Forthcoming in “Guide to the Philosophy of Cosmology”, Ijjas and Loewer (ed.) Preprint at <http://philsci-archive.pitt.edu/15295/>.
- Wallace, D. (2016b). Probability and irreversibility in modern statistical mechanics: Classical and quantum. To appear in D. Bedingham, O. Maroney and C. Timpson (eds.), *Quantum Foundations of Statistical Mechanics* (Oxford University Press, forthcoming). Preprint at <https://arxiv.org/abs/2104.11223>.
- Wallace, D. (2019a). Isolated systems and their symmetries, part II: Local and global symmetries of field theories. Forthcoming in *Studies in the History and Philosophy of Science*. Preprint at <http://philsci-archive.pitt.edu/19729/>.
- Wallace, D. (2019b). Naturalness and emergence. *Monist* 102, 499–524.
- Wallace, D. (2019c). What is orthodox quantum mechanics? In A. Cordero (Ed.), *Philosophers Look at Quantum Mechanics*, pp. 285–309. Switzerland: Springer.
- Weinberg, S. (1979). Ultraviolet divergences in quantum theories of gravitation. In S. Hawking and W. Israel (Eds.), *General Relativity: an Einstein Centenary Survey*, pp. 790–832. Cambridge: Cambridge University Press.
- Weinberg, S. (1989). The cosmological constant problem. *Reviews of Modern Physics* 61, 1–23.
- Weinberg, S. (1995). *The Quantum Theory of Fields, Volume II: Modern Applications*. Cambridge: Cambridge University Press.
- Weinberg, S. (2008). *Cosmology*. Oxford: Oxford University Press.
- Wiltshire, D. L. (2011). What is dust? - physics foundations of the averaging problem in cosmology. *Classical and Quantum Gravity* 28, 164006.
- Zurek, W. H. and J. P. Paz (1999). Why we don’t need quantum planetary dynamics: Decoherence and the correspondence principle for chaotic systems. In D. Greenberger, W. L. Reiter, and A. Zeilinger (Eds.), *Epistemological and Experimental Perspectives on Quantum Physics*. Dordrecht: Springer. <https://arxiv.org/abs/quant-ph/9612037>.