UNIVERSITY OF COPENHAGEN

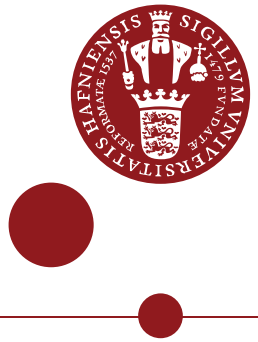

MSc in Quantum Physics

Black Hole Mimickers from Asymptotically Safe Einstein-Weyl Gravity

Jonas Pfeiffer Svarre Rasmussen

Supervised by Alessia Platania

20th May 2025

Jonas Pfeiffer Svarre Rasmussen
Black Hole Mimickers from Asymptotically Safe Einstein-Weyl Gravity
MSc in Quantum Physics, May 20th, 2025
Supervisors: Alessia Platania (UCPH)

University of Copenhagen
Faculty of Science
Niels Bohr Institute
Dep.
Blegdamsvej 17
2100 Copenhagen Ø, Denmark

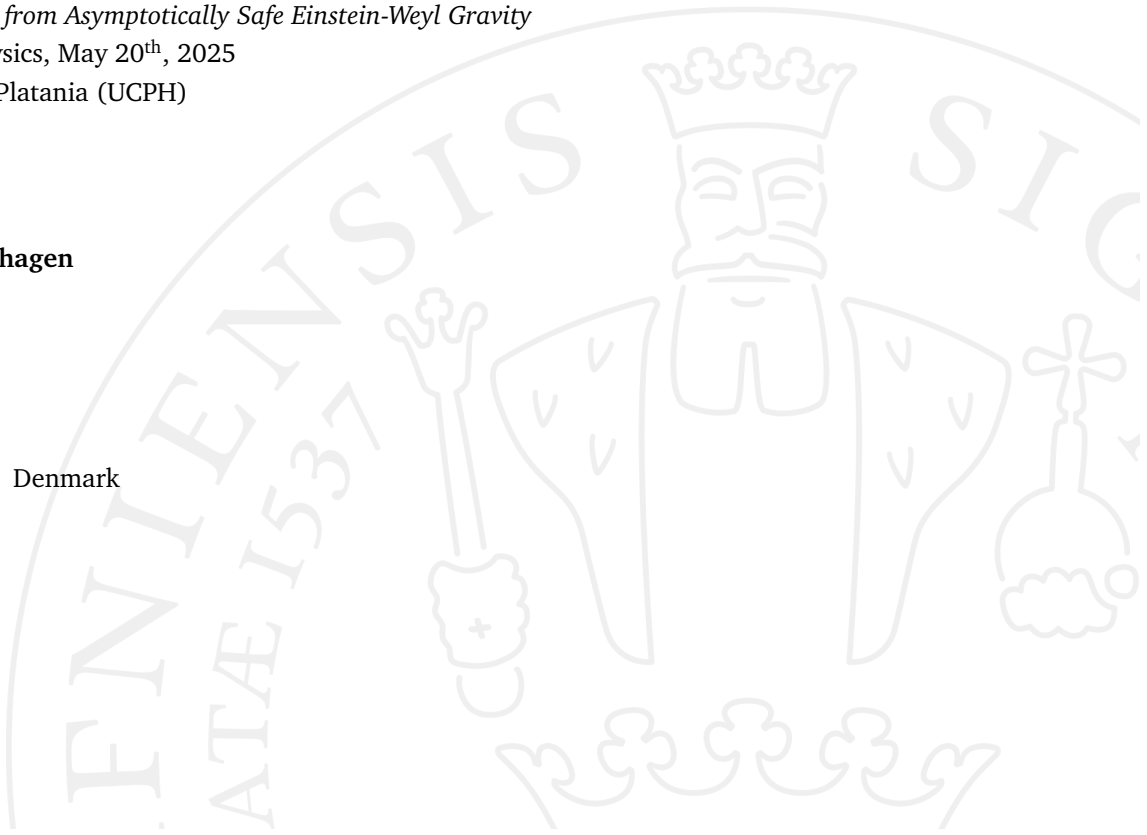

Abstract

This thesis explores the interface between quantum gravity and black hole physics within the framework of asymptotic safety. The asymptotic safety scenario posits that a quantum field theory of gravity may be ultraviolet complete if its renormalization group flow is attracted to an interacting ultraviolet completion at high energies. Using functional renormalization group techniques applied to a truncated gravitational action, we provide a proof of principle of how the landscape of black hole-like spacetimes can be mapped out from fundamental physics.

We focus on the Einstein-Weyl approximation — the simplest extension of general relativity, where the Einstein–Hilbert term is complemented by the Weyl-square invariant — and we compute the effective action stemming from its corresponding asymptotically safe ultraviolet completion. We use a first-principle calculation based on beta functions, which encode the renormalization group flow of the theory, and determine the unique ultraviolet-complete trajectory connecting this fixed point to the infrared. Along this trajectory, the Wilson coefficient of the Weyl-squared term is fixed to $G_{\mathrm{C}^2} = 0.5092\, m_{\mathrm{Pl}}^{-2}$.

We analyse the implications of the resulting effective field theory for the classically derived phase diagram of black hole mimickers in Einstein–Weyl gravity. This diagram maps the different solution families — including Schwarzschild and non-Schwarzschild black holes, asymmetric wormholes, and naked singularities — as a function of the object's gravitational parameters, such as its mass M and the Wilson coefficient G_{C^2}. Inputting the quantum gravity-determined Wilson coefficient for the Weyl-squared term constrains the phase diagram and the set of allowed solutions. We find in particular that asymptotic safety disfavours Bachian naked singularities.

These results provide a concrete example of how quantum gravity can yield constraints on black hole physics and offer a proof of concept for the strength and predictive power of the asymptotic safety framework.

Contents

1 Introduction

The quest for a consistent and predictive theory of quantum gravity remains one of the great open challenges in fundamental physics. While General Relativity (GR) continues to match observations across the solar system and astrophysical scales with remarkable precision, it is also clear that it must eventually give way to a quantum description.

This is not just because singularities — such as those inside of black holes or at the Big Bang — signal a breakdown of the theory, but because gravity is believed to be a fundamental interaction alongside the strong and electroweak interactions; these are known to be quantum and are described within the Standard Model (SM) of particle physics.

From a theoretical perspective, the problem is clear: When treated as a quantum field theory (QFT), GR is not perturbatively renormalizable. Because the Newton's coupling has a negative mass dimension, new divergences appear at each loop order, requiring an infinite number of counterterms and rendering the theory non-predictive in the ultraviolet (UV). Over the years, many proposals have aimed to address this issue, ranging from string theory to loop quantum gravity. This will all be discussed in Chapter 2. A more conservative approach — and the one we pursue in this thesis — is known as asymptotically safe quantum gravity (ASQG). Originally proposed by Weinberg in the 1970s [1], this scenario posits that gravity might be well-defined at all energy scales and non-perturbatively renormalizable if its renormalization group (RG) flow approaches a non-Gaussian — i.e. interacting — fixed point (NGFP) at high energies.

In recent years, evidence for the existence of such a fixed point has steadily accumulated, particularly through the use of the functional renormalization group (FRG) [2, 3, 4]. The FRG offers a powerful non-perturbative tool to study RG flows in theory space, based on a scale-dependent effective action that evolves according to the Wetterich equation. The earliest studies focused on the Einstein-Hilbert truncation [5, 6, 7] — involving only Newton's constant and the cosmological constant — and later works have built upon this work by including matter fields and higher curvature terms [8, 9, 10, 11]. These efforts point toward a consistent picture: a finite-dimensional critical

surface in theory space and the tantalising possibility of a UV-complete quantum theory of gravity.

But the fixed point itself is only part of the story. If asymptotic safety is to describe the real world, it must leave an imprint at macroscopic scales — in observables that are, at least in principle, accessible. Black holes provide a natural testing ground. Their strong gravitational fields — which are unlike anything we can hope to probe at Earth — and their thermodynamic properties touch on quantum effects in curved spacetime, and their internal structure — often hidden behind horizons — remains a mystery. Classical GR predicts singularities inside black holes, which are — in the truest sense of the word — unphysical; a quantum theory of gravity should resolve or reinterpret them.

Independent of the specific approach to quantum gravity, researchers have tried to incorporate, or at least parametrise, quantum corrections to classical black hole spacetimes. Some modify the metric by hand, others use semiclassical arguments or scale-dependent couplings inspired by RG flow. While these methods yield interesting insights, also in the context of asymptotic safety [12], they often rely on heuristic choices or specific model-building assumptions [13]. What has been missing is a more systematic approach — one that derives modified black hole solutions from first principles and, in particular, from a quantum effective action stemming from a given UV completion — for instance, an asymptotically safe one. This thesis provides a proof of principle for such a systematic study.

A key part of this thesis is devoted to building physical intuition for the underpinning concepts of asymptotic safety — in particular, effective field theory (EFT), theory space, and RG flow, which we will explore in Chapter 3. Starting from the Ising model and coarse-graining procedures in statistical physics, we develop the conceptual machinery that connects microscopic details to large-scale behaviour. This journey not only clarifies why the framework of effective actions is a fundamental tool in theoretical physics but also provides the natural language for exploring how gravity might flow to a fixed point in the UV. These discussions are essential for interpreting the results that follow — and for appreciating how something as abstract as the structure of spacetime can be built from a language originally intended to describe something as seemingly simple as water turning to ice [14].

After these preliminary chapters we will review the framework of ASQG in Chapter 4, highlighting its key achievements and drawbacks, and introducing the concept of "asymptotic safety landscape" which will be crucial for the scope of this thesis: constructing the "phase diagram" of all possible black hole-like solutions stemming from asymptotic safety in the so-called Einstein-Weyl truncation. This truncation corrects

GR by the inclusion of the Weyl-squared term. The solutions to the resulting field equations span a rich landscape. The classically derived phase diagram of black hole mimickers — characterised by the asymptotic mass and Yukawa charge — includes Schwarzschild and non-Schwarzschild black holes, wormholes, and other exotic compact objects [15], as reviewed in Chapter 5. However, the landscape of these solutions depends on the Wilson coefficient of the Weyl-squared term. The key idea underlying this thesis and developed in Chapter 6 is that this coupling ought to be fixed by quantum gravity, and in particular, in ASQG, it can be feasibly computed.

Using beta functions derived elsewhere [8] from the FRG, in Chapter 6 we study the RG flow in the Einstein-Weyl truncation. Crucially, we show that the requirement of asymptotic safety selects a single, UV-complete RG trajectory that connects the NGFP to the infrared (IR). Along this trajectory, the Wilson coefficient of the Weyl-squared term is uniquely fixed. In other words: if asymptotic safety holds in this truncation, the effective action no longer contains free parameters — it is uniquely determined by the fixed point, namely, the UV-completion. After isolating the unique UV-complete trajectory and following it to the IR, we compute the corresponding Wilson coefficient. This thus places a constraint on which solutions can be physically realised in a scenario where gravity is asymptotically safe. Among other things, we find that Bachian naked singularities are disfavoured in asymptotic safety. These results are derived in Chapter 6.

While this thesis does not claim to present the final word on quantum black holes, it offers a proof of principle: a concrete way to connect the microscopic dynamics of asymptotic safety to the macroscopic structure of spacetime. It will serve as an example of how quantum gravity phenomenology — normally done via model building and classical reasoning — can be guided by fundamental physics. Discussion on possible drawbacks and extensions of our work are reported in Chapter 7 and 8.

2 Towards Quantum Gravity

2.1 Why Quantum Gravity?

The two most successful theories of modern physics — SM and GR — describe the part of nature they individually govern with extraordinary precision. But they speak fundamentally different languages: one describes the quantized matter and force-carrying particles of the SM using QFT on a background of flat spacetime, while the other treats spacetime itself as a classical, dynamical background where the matter fields back-react dynamically with the background. The actors form the spacetime — and the spacetime tells the actors how to act gravitationally. As a holistic argument, this seems inconsistent — why would the "stage" the quantum "actors" act on, not itself be quantum?

This divide is more than skin-deep. In any system where quantum matter significantly affects the curvature of spacetime — such as in the early universe or near black hole singularities — a hybrid framework of the two descriptions breaks down. A consistent description of nature must treat both the stage and the actors quantum mechanically. Semi-classical approaches like $G_{\mu\nu} = \langle T_{\mu\nu} \rangle$ are useful approximations, but they fail when quantum fluctuations are large.

And then there are the singularities. GR predicts that spacetime simply *ends* inside a black hole. Singularities are — in the truest sense of the word — unphysical, and signal the failure of a theory at that specific point. Just as quantum mechanics resolved the UV catastrophe of classical electromagnetism, we would expect a quantum theory of gravity to resolve — or at least reinterpret — gravitational singularities.

All of this hints at the quantum nature of gravity. However, as we shall see in the following, when treated as a QFT, GR turns out to be perturbatively non-renormalizable. This means that the theory develops divergences that cannot be reabsorbed in a finite number of parameters. A quantum theory of gravity thus has to be developed, extending in some direction the well-established framework of perturbative QFT, which is utilised to describe the other fundamental interactions of nature.

2.2 Perturbative Non-Renormalizability of General Relativity

GR, as encoded in the Einstein-Hilbert action, faces a glaring problem when treated as a QFT: it is perturbatively non-renormalizable. This is not just a technical annoyance — it indicates that GR cannot be the "complete story" at high energies. Let us unpack what this means, first mathematically, then physically.

The Einstein-Hilbert action is, in standard Lorentzian signature, given by:

$$S_{\text{EH}} = \frac{1}{16\pi G_{\text{N}}} \int d^4x \sqrt{-g}\, R\,, \tag{2.1}$$

where G_{N} is Newton's constant, g is the determinant of the metric $g = \det(g_{\mu\nu})$ and R is the Ricci scalar, which importantly contains second-order derivatives of the metric. To quantize gravity perturbatively, one usually expands the metric around flat spacetime:

$$g_{\mu\nu} = \eta_{\mu\nu} + \kappa h_{\mu\nu}\,, \quad \text{with } \kappa = \sqrt{32\pi G_{\text{N}}}\,, \tag{2.2}$$

where $\eta_{\mu\nu}$ is the Minkowski metric, and the field fluctuating on the flat background $h_{\mu\nu}$ represents the graviton. This expansion yields an infinite number of interaction vertices due to the nonlinear structure of the action $\sqrt{-g}R$.

The graviton propagator, derived as the "inverse" of the quadratic part of the action, behaves like:

$$\Delta(p) \sim \frac{1}{p^2}\,, \tag{2.3}$$

which is similar to the massless photon in Quantum Electrodynamics (QED). However, the crucial difference lies in the dimension of the coupling constant.

In four dimensions, Newton's constant has mass dimension:

$$[G_{\text{N}}] = -2\,. \tag{2.4}$$

This implies that the coupling κ has mass dimension $[\kappa] = -1$. In QFT, this signals that the theory becomes strongly coupled at high energies. This is not a problem for on-shell, tree-level Feynman diagrams. However, loop diagrams introduce integrations over virtual momenta, which tear things apart. To see the problem, consider a generic loop diagram. As in standard textbooks on QFT, we can, using power counting of the

mass dimensions, find the *superficial degree of divergence* D. For a diagram with L loops, I internal lines, and V vertices, D (in 4 spacetime dimensions) is given by

$$D \sim 4L - 2I + 2V \,, \tag{2.5}$$

where:

- $4L$ comes from the loop momentum integrals $\int d^4k$,
- $-2I$ accounts for the graviton propagators which go like $\sim 1/k^2$,
- $+2V$ reflects the fact that each graviton vertex introduces two derivatives, i.e., two powers of momentum, due to second-order derivatives of the metric in the Ricci scalar.

A diagram is (usually) UV divergent if $D \geq 0$, and convergent if $D < 0$ [16]. Using the topological identity $L = I - V + 1$, we can easily see that

$$D = 2L + 2 \,, \tag{2.6}$$

showing explicitly that divergences worsen at higher loops.

In standard perturbative QFT, one can renormalize by introducing *counterterms* to counteract these divergences, such that "one infinity swallows another infinity", leaving us with something finite. In exchange for finiteness, the price is one more parameter for each counterterm, and hence, a new observable that must be measured. As an example, QED is a renormalizable quantum field theory in which the superficial degree of divergence decreases with loop order. It has only three divergent diagrams, each of which can be cancelled by introducing a corresponding counterterm. These correspond to three physical input parameters: the mass and charge of the electron are the two recognisable ones.[1]

In contrast, because the gravitational coupling κ has negative mass dimension and every graviton vertex introduces momentum dependence, an infinite number of counterterms are required to absorb all UV divergences in GR. This renders the theory perturbatively non-renormalizable: it loses its predictive power at high energies, such as near the centre of a black hole, where quantum gravitational effects become significant.

[1] Photon field strength normalisation is the third one.

And although we happily use GR in low-energy physics for planetary motion, gravitational waves, cosmology and time dilation in satellites, as it is an excellent EFT[2], its breakdown in the UV makes it clear: **GR is not the end of the story, and its perturbatively quantized version cannot be a fundamental quantum theory of gravity**.

2.3 Beyond Quantum GR: Approaches to Quantum Gravity

The non-renormalizability of GR, combined with the slew of other reasons for a quantum theory of gravity, has compelled physicists to search for a UV-complete theory of quantum gravity. Several approaches attempt this [4, 17], including:

- **String theory**, where gravity emerges from quantized strings, and UV divergences are softened by the extended nature of strings.
- **Loop quantum gravity**, which quantizes the geometry of spacetime using holonomies and weak constraints.
- **ASQG**, which retains the QFT framework based on a non-perturbative UV completion given by an interacting fixed point of the RG flow.

In the following chapters, we will focus on and explore the last of these — asymptotic safety. The reason is that the latter is the most conservative approach to quantum gravity, and its formulation, which retains a QFT formalism, makes it feasible to provide a proof of principle of how quantum black holes may be derived or constrained by the requirement of UV completion. To explore these concepts and arrive at the results of this thesis, we shall first introduce the Wilsonian RG and the concept of asymptotic safety in quantum gravity.

[2] More on this in the next two chapters.

3 The Wilsonian Picture of Renormalization

It is a counterintuitive fact that the physics of the large is often far simpler than the physics of the small. Water freezes at $0°$C, iron becomes magnetised below a critical temperature, and gases obey simple thermodynamic laws. It seems almost miraculous, as these are all emergent laws, rising out of the mind-numbing complexity of the microscopic. A glass of water contains roughly 10^{24} molecules. Each molecule is made of atoms, whose electrons move around dense, quantum-fluctuating nuclei made of protons and neutrons. In turn, these are bound states of quarks and gluons governed by quantum chromodynamics. Dig even deeper, and you might find vibrating strings or even more speculative structures beneath.

So, how can we even begin to understand the first thing about water, let alone describe it with a few simple equations, without first solving all the deep mysteries of beyond-standard-model physics and quantum gravity? The answer lies in the concept of **renormalization**.

Although many textbooks would have you believe otherwise, renormalization is not just a trick for taming infinities in QFT. It is, more fundamentally, the *microscope* of theoretical physics. It provides a framework for understanding how microscopic complexity gives rise to macroscopic simplicity; it tells us which details matter and which can be safely ignored as we zoom out.

Statistical physics teaches us that concepts like temperature and pressure are not properties of individual particles — there is no sense in defining the pressure of a single atom floating about — they emerge from averaging over ensembles. Renormalization tells us how those averages behave as we change the scale at which we observe a system. It also explains why many microscopic systems exhibit the same macroscopic behaviour near critical points, a phenomenon known as *universality*.

Fueled by the idea of asking the right questions at the right scale, in this chapter, we lay the conceptual foundation for what follows. We begin with what seems like a toy model — the Ising model — to develop the central ideas behind coarse-graining and scale dependence. From there, we generalise to continuous fields and introduce the Wilsonian view of renormalization. We will then put these concepts together

and use them to introduce the framework of asymptotic safety in quantum gravity in Chapter 4.

3.1 The Ising Model and Zooming Out via Coarse-Graining

Let us begin with a classic physics problem: the Ising model. While it does not immediately bear any relevance for our system, its behaviour will be an intuitive starting point for discussing the Wilsonian picture of renormalization.
Imagine two particles with spin in a magnetic field. The energy of this simple system can be described by the Hamiltonian:

$$H = -B(s_1 + s_2) - J \cdot s_1 s_2 \,, \tag{3.1}$$

where B is some external magnetic field, and J is an interaction constant. Throughout this thesis, we will use *actions*, not Hamiltonians, which we use here for demonstration purposes. However, the role of this Hamiltonian will be identical to the role of the Lagrangians we will encounter.
We can generalise this system to a macroscopic system:

$$H = -B \sum_{i}^{N} s_i - J \sum_{\langle ij \rangle}^{N} s_i s_j \,, \tag{3.2}$$

where we have made the simple assumption that particles only interact with their immediate neighbour, denoted by the notation $\langle ij \rangle$. In standard statistical mechanics fashion, we can calculate the probability of being in a specific configuration by using the Boltzmann distribution

$$p(s) = \frac{e^{-\beta H(s)}}{Z} \,, \qquad\qquad \mathcal{Z} = \sum_{s} e^{-\beta H(s)} \,. \tag{3.3}$$

Here $\mathcal{Z}$ is the partition function and β is the usual inverse temperature times the inverse Boltzmann constant $\beta = (k_{\mathrm{B}} T)^{-1}$. We sum over $s = \{s_i\}$, which is a specific configuration of the spins.

In principle, we now know all there is to know about our system — H governs the microscopic description, and "induces" all the large-scale physics in $\mathcal{Z}$. Indeed, the latter can give us information on all the macroscopic thermodynamic quantities: pressure, average energy, magnetisation and more.

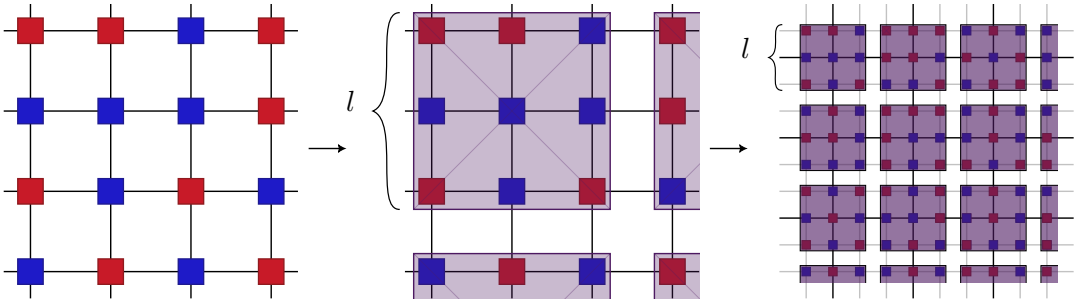

FIGURE 3.1.
The procedure of coarse-graining. *Left panel*: a square lattice of spin-up (red) and spin-down (blue) atoms. *Middle panel*: lattice sites grouped into $l \times l$ blocks. *Right panel*: each coarse-grained block represented by a local magnetization $m(x) \in [-1, 1]$.

There is a problem, however — it is almost impossible in practice.

Consider computing the average magnetisation of our system:

$$\langle m \rangle = \frac{1}{N} \langle \sum_i^N s_i \rangle = \frac{1}{N} \sum_s p(s) \sum_i^N s_i = \frac{1}{N} \sum_s \frac{e^{-\beta H(s_i)}}{\mathcal{Z}} \sum_i^N s_i = \frac{1}{N\beta} \frac{\partial \log \mathcal{Z}}{\partial B} , \quad (3.4)$$

where $m = \frac{1}{N} \sum_i^N s_i$ is the magnetisation of a specific configuration of spins $s = \{s_i\}$. The issue is that the number of possible spin configurations grows exponentially with N, i.e. 2^N. This means that obtaining p or $\mathcal{Z}$ for any macroscopic object (say about $\sim 10^{23}$ particles) would be a tedious task to put it mildly. To access the large-scale properties hidden in $\mathcal{Z}$ behind the enormous complexity of the microscopic description, we need some kind of **in-between** description that connects microscopic interactions to macroscopic behaviour.

3.1.1 Coarse-Graining

The following derivation is intentionally crude — our goal is to grasp the intuitive idea of coarse-graining and not get lost in mathematical formalism.

Imagine a 2D square lattice of spins arranged with uniform spacing. We can divide the system into blocks of side length l, thus encapsulating $l \times l = n$ spins in each. For each block, we define the average/*coarse-grained* magnetisation:

$$m(x) = \frac{1}{l^2} \sum_i^n s_i , \quad (3.5)$$

where x denotes the centre of the block. An example of this procedure with a 3×3 block can be seen in Figure 3.1.

Although both m and x are technically discrete (since the system has a finite number of particles with finite spacing), on macroscopic scales of something like a table made

from iron, we can safely treat them as continuous. This means that we have effectively promoted $m(x)$ to a continuous magnetisation field defined over position x. This leads us to reformulate the p and $\mathcal{Z}$ in a continuous language:

$$p(s) \to p_{eff}[m(x)] = \frac{\mathrm{e}^{-\beta H_{eff}[m(x)]}}{\mathcal{Z}}, \qquad \mathcal{Z} \to \mathcal{Z} = \int \mathcal{D}m \mathrm{e}^{-\beta H_{eff}[m(x)]}. \tag{3.6}$$

Here, $H_{eff}[m(x)]$ is some effective Hamiltonian that describes the coarse-grained degrees of freedom,[1] the probability $p[m(x)]$ is a functional of the field $m(x)$ and the partition function $\mathcal{Z}$ is now a functional integral or *path integral* (which justifies the introduction of the measure $\mathcal{D}m$).

This is a major conceptual leap. We now deal with far fewer degrees of freedom. For example, with 3×3 blocks, each coarse-grained spin site has one spin value $m \in [-1, 1]$ instead of 9, in principle making our job of calculating the average magnetisation 2^9 times easier! A crucial point is that the partition function remains unchanged (hence the lack of any subscript). We require it to encode the same macroscopic physics — pressure, magnetisation, etc. — i.e., coarse-graining should not alter what we measure at large scales; it rewrites the microscopic theory in a more convenient form.

We should stress at this point that we are transforming the *description*, not the *physics*. If coarse-graining changed the value of $\mathcal{Z}$, then observables like pressure or average energy would also change, which would be unphysical. So no matter how we slice and dice our theory, the partition function should remain unchanged $\mathcal{Z} \to \mathcal{Z}$.

Of course, we are not limited to just one round of coarse-graining. The real power of this procedure emerges when we iterate it: after coarse-graining into $l \times l$ blocks, we can repeat the process by grouping the new blocks (which now function as a single lattice site) into even larger $l \times l$ blocks, defining a new, smoother field at each step. By repeatedly "zooming out" in this way, we systematically strip away microscopic detail, flowing toward a description that captures only the long-wavelength, large-scale behaviour. This is the essence of the RG: understanding how the effective description of a system evolves as we successively change the scale at which we view it.

3.1.2 Splitting Up the Field

We now build a more robust description of coarse-graining by describing it as a decomposition of our magnetisation field into low- and high-energy components.

[1] This Hamiltonian is analogous to a thermodynamic free energy [14].

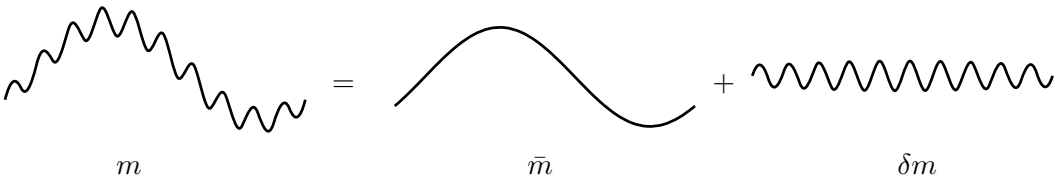

FIGURE 3.2.
Splitting the field into smooth background $\bar{m}(x)$ and fluctuations $\delta m(x)$.

Imagine you are hovering in an aircraft just a couple of meters above the Pacific Ocean. To pass the time, you focus your attention on the tidal motion of the surface: massive swells threatening to engulf your ship, gentle mid-scale waves, and tiny ripples scattered across the surface. As a sensing being — organic or robotic — you are limited by some resolution, whichever way you want to perceive the world around you. Just like a camera cannot take a picture with infinite pixels, you cannot see a speck of dust on the moon or hear somebody whispering in Hawaii;[2] nothing has infinite resolution.

Sure enough, as you turn on your aircraft and ascend, detail is progressively lost. From 10 meters above, the tiny ripples vanish. At 100 meters, the mid-scale waves blur. At 1 km, only the deep blue remains. Thus, zooming out makes fine details disappear. This is exactly what goes on in the RG equations. Just as the ocean's surface is a messy superposition of ripples and waves, our magnetisation field $m(x)$ is a blend of fluctuations across scales. We can formalise this by splitting $m(x)$ into the sum of a background field and a fluctuation field:

$$m(x) = \bar{m}(x) + \delta m(x)\,, \tag{3.7}$$

with some momentum scale k determining the split's cutoff, as illustrated in Figure 3.2. In the language of the ocean analogy, the background field $\bar{m}$ contains the long-wavelength modes (waves) which are still detectable, whereas δm encodes the short-wavelength fluctuations (ripples) with fine structure undetectable at our current resolution. Instead of simply stating that we have some "effective" Hamiltonian as before, let us now formulate a real one.

By demanding locality, continuous translational and rotational invariance,[3] analyticity, and setting the external magnetic field to $B = 0$ so that the Hamiltonian respects a $\mathbb{Z}_2$ symmetry $m(x) \to -m(x)$, we can write down the most general local Hamiltonian as an expansion in even powers of the magnetisation field $m(x)$:

$$H[m(x)] = \int d^d x \left(g_2 m^2(x) + g_4 m^4(x) + b_2 (\nabla m(x))^2 + \dots \right)\,, \tag{3.8}$$

[2] Under the assumption that you are neither on the moon nor in Hawaii.
[3] If we clenched to the discreteness of our system, it would only be invariant under discrete versions of these symmetry transformations.

where all odd powers of $m(x)$ must go to zero due to the $\mathbb{Z}_2$-symmetry, and the coefficients $\{g_i, b_i\}$ generally depend on external parameters such as temperature, external magnetic field B and, importantly, a *renormalization scale* k at which we coarse-grain. The gradient term encodes spatial stiffness, which penalises sharp fluctuations in $m(x)$.

This form of the Hamiltonian, known as the Ginzburg–Landau Hamiltonian, is a classic textbook example which models systems near criticality, such as ferromagnets, superfluids, and superconductors [14].

We can now coarse-grain more formally by *integrating out* the high-momentum fluctuations $\delta m(x)$. Doing so modifies the Hamiltonian H and the effective couplings $\{g_i, b_i\}$, which change with the cutoff scale k. More precisely, when $m(x) \to \bar{m}(x)$, then $H[m(x)] \to \bar{H}[\bar{m}(x)]$ and $\{g_i, b_i\} \to \{\bar{g}_i, \bar{b}_i\}$. To track the inner workings of our system, we focus on how the probability of a given field configuration changes with scale. Specifically, what is the probability of observing a coarse-grained field $\bar{m}(x)$, and how does that relate to the original probability distribution over $m(x)$? We require the functional form of the distribution to remain unchanged under coarse-graining:

$$p[m(x)] = \frac{e^{-H}}{\mathcal{Z}} \xrightarrow{m \to \bar{m}} p[\bar{m}(x)] = \frac{e^{-\bar{H}}}{\mathcal{Z}}\,, \tag{3.9}$$

where we have omitted β, as it is not important for our calculations, and as before, the partition function is kept unchanged under this procedure:

$$\mathcal{Z} = \int \mathcal{D}m\, e^{-\beta H[m(x)]} = \int \mathcal{D}\bar{m}\, e^{-\beta \bar{H}[\bar{m}(x)]}\,. \tag{3.10}$$

Without delving too far into mathematical rigour, we can phrase it as such: The probability of ending up in a specific coarse-grained configuration $\bar{m}$ must be equal to the total probability of all fine-grained configurations m that average to $\bar{m}$ when we coarse-grain. Formally,

$$p[\bar{m}] \sim \int_{m \to \bar{m}} \mathcal{D}m \mathrm{e}^{-H[m]} = \int_{\delta m} \mathcal{D}m \mathrm{e}^{-H[\bar{m}+\delta m]} = \mathrm{e}^{-\bar{H}[\bar{m}]}\,, \tag{3.11}$$

where $\bar{H}[\bar{m}]$ is the renormalized Hamiltonian describing the effective theory at scale k. We can expand the Hamiltonian before and after integrating out δm to illustrate the structure:

$$\begin{aligned} H[m] &= H[\bar{m} + \delta m] \\ &= \int d^d x \left(g_2(\bar{m} + \delta m)^2 + g_4(\bar{m} + \delta m)^4 + b_2(\nabla(\bar{m} + \delta m))^2 + \dots \right)\,. \end{aligned} \tag{3.12}$$

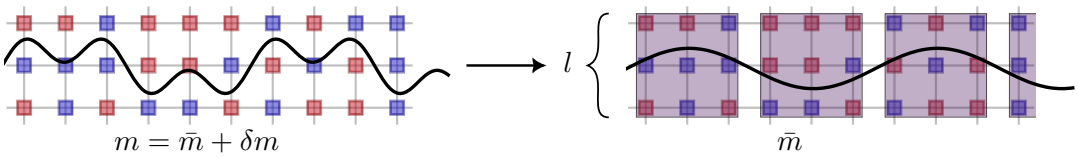

FIGURE 3.3.
Decomposing the field and coarse-graining: small-scale details are averaged out, redefining both $\bar{m}(x)$ and the couplings.

Doing this calculation, rearranging and matching terms of the same order of $\bar{m}(x)$ such that $H[m]$ keep the same structure as $\bar{H}[\bar{m}]$, we can redefine terms containing non-trivial combinations of $\{g_i, b_i\}$ and δm to new couplings constants $\{\bar{g}_i, \bar{b}_i\}$. So, introducing a cutoff scale changes the couplings $\{g_i, b_i\} \to \{\bar{g}_i, \bar{b}_i\}$ where the new renormalized couplings are functions of $(g_i, b_i, B, T, \delta m)$,

$$H[m + \delta m] \xrightarrow{\text{integrate out } \delta m} \bar{H}[\bar{m}] = \int d^d x \big(\bar{g}_2 \bar{m}^2 + \bar{g}_4 \bar{m}^4 + \bar{b}_2 (\nabla \bar{m})^2 + \mathcal{O}(g_i, b_i, \delta m, \nabla \delta m, \bar{m}^6, \dots) + \dots \big) . \tag{3.13}$$

In essence — as illustrated in Figure 3.3 — changing the cutoff frequency, which separates "slow" and "fast" modes, is completely analogous to coarse-graining over an $l \times l$ block — rising above the ocean. In all cases, small-scale fluctuations are smoothed out, and the effective description evolves. As emphasised earlier, there is no reason to stop at a single round of coarse-graining. The power of the formalism we have just introduced — splitting the field into low- and high-momentum components and integrating out the high-momentum modes — is that it naturally allows (almost begs!) to be iterated. This allows us to smoothly zoom out little by little, peeling away layers of microscopic detail, until we find ourselves in the macroscopic world, where rich and often universal physics begins to emerge. We will return to discuss a truncated toy model such as this in the next section.

3.2 Wilsonian Picture of Renormalization

So far, the Hamiltonians we have encountered have been described from a pragmatic, "bottom-up" perspective: we write down the most general local action (or Hamiltonian) allowed by symmetries and truncate by relevance. But this leaves (at least) two questions unanswered: Why does this expansion structure emerge at all? And how does the theory know which operators to keep and which to suppress?

These questions are tackled elegantly by the Wilsonian picture of renormalization, as encoded, for instance, in the Wilson-Polchinski renormalization group equation (RGE) [18]. The latter treats the scale dependence of the effective action in a dynamical fashion, by successively integrating out momentum shells — just as we

have seen in the toy model example when we split up the fields in slow and fast modes. Using this method, we can define a flow on the space of all possible actions, which gives a precise mechanism for the decoupling of higher-order terms. This can also be done in a functional fashion, resulting in an exact RG equation alternative to the Wilson-Polchinski one, which is called the Wetterich equation [19] and will be introduced at the end of this Chapter.

3.2.1 Exploring Theory Space: Toy Model

In the last section, we introduced the intuitive notion of coarse-graining via the Ising model. In this section, we build on that foundation by walking through a concrete, worked example that brings several abstract ideas to life, which are key ingredients of the Wilsonian RG. Our goal is to explore how a theory evolves as we progressively integrate out high-energy degrees of freedom, giving a intuitive understanding of the core concepts of RG flow, theory space, and fixed points along the way. Fixed points in particular play a crucial role in relation to renormalization, as we shall see in the next subsections.

We start by setting k to some high, but finite, arbitrary cutoff $k = \Lambda$. In the context of QFT, a theory is well-defined if all its couplings remain finite in the limit $\Lambda \to \infty$ — a point which will become clear shortly — but for the purpose of this section, we will simply start from a finite $k = \Lambda$. From there, we integrate out the high-momentum fluctuations δm, leading to the transformation:

$$\{H, m, g_i, b_i\} \longrightarrow \{\bar{H}, \bar{m}, \bar{g}_i, \bar{b}_i\}\,.$$

Having set our high-scale cutoff, we now lower it slightly, $k \to \Lambda - \delta$, thereby splitting the current background field $\bar{m}$ once again into a new set of slow modes and fast fluctuations:

$$\bar{m}(x) = \bar{\bar{m}}(x) + \delta\bar{m}(x)\,. \tag{3.14}$$

Wash, rinse, repeat — each step smooths out a little more detail, progressively revealing the long-distance physics. This recursive coarse-graining is the core mechanism behind what makes the theory *flow* to different values at different scales — appropriately dubbed the *RG flow* of the theory. We are gonna explore the intuition behind this using an arbitrary Hamiltonian truncated to two terms:

$$H[m] = \int d^dx \left(\sum_{n=even} [g_n m^n(x) + b_n(\nabla m)^n] \right) \to \int d^dx \left((g_2 m^2(x) + g_4 m^4(x)) \right)\,.$$

Thus, we are focusing on a 2-dimensional subset of *theory space*. Theory space, simply put, is the (often infinite-dimensional) coordinate space whose axes are spanned by

the couplings multiplying every operator consistent with the spacetime dimension, field content, regulator, and symmetries, such that each point specifies a distinct QFT. By truncating up to a certain order, we thus consider an n-dimensional subset of theory space, where n is the number of terms left in the Hamiltonian (or Lagrangian/action).

We choose an initial theory by fixing the value for the two couplings $\{g_2, g_4\}$. This, we can depict as a point in theory space, like the one in the left panel of Figure 3.4. Now that we are seasoned in the process of splitting up the field and redefining our system, we can do it in one fell swoop:

$$H[m] = H[\bar{m} + \delta m] = \int d^d x \left(g_2(\bar{m} + \delta m)^2 + g_4(\bar{m} + \delta m)^4 \right) \tag{3.15}$$

$$\xrightarrow{\text{integrate out } \delta m} H[\bar{m}] = \int d^d x \left(\bar{g}_2 \bar{m}^2 + \bar{g}_4 \bar{m}^4 \right) , \tag{3.16}$$

where we omit the higher-order crossing terms between the different fields to keep the truncation consistent.[4] This process moves our theory to a new position in theory space $\{\bar{g}_2, \bar{g}_4\}$ — the couplings *runs* with cutoff scale k, so to speak. This is depicted in the middle panel of Figure 3.4. If we keep reiterating the process of coarse-graining and redefining (i.e., zooming more and more out by letting $k \to 0$), we have two outcomes:

1. The couplings of our theory keep growing in magnitude and go to infinity.

2. The couplings approach a specific value — we end up at a *fixed point*.

The second possibility is the one of physical interest, which will be backed up in the context of fixed points and predictability in Section 3.2.3.

The path the theory took in theory space as we renormalized it is called the *trajectory* of the theory. If we imagine testing a slew of different theories with the same truncation — i.e. same coupling terms $\{g_2, g_4\}$ — but with different initial values, and mapping their trajectories through theory space, a curious picture emerges. A *flow diagram*, which shows how the theory *flows* through theory space as the *RG scale* k (i.e., the cutoff frequency) is varied. The RG flow is exactly described by the coupled system of *beta functions*, where each arrow represents the vector

$$\vec{\beta}(g_2, g_4) = \begin{bmatrix} \beta_2(g_2, g_4) \\ \beta_4(g_2, g_4) \end{bmatrix} = \begin{bmatrix} k\partial_k g_2 \\ k\partial_k g_4 \end{bmatrix} . \tag{3.17}$$

[4] In reality, this is a much more involved process where we use the requirement of keeping the probability on the same form like in (3.9) and then expanding in powers of δm.

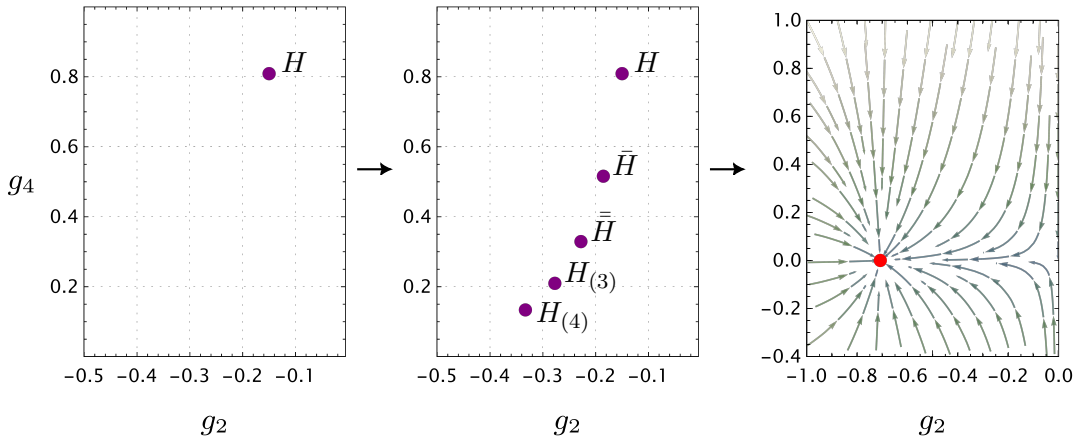

FIGURE 3.4.
Left panel: Some initial theory in a two-dimensional theory space. *Middle panel*: How coarse-graining moves the theory in theory space. *Right panel*: Flow diagram of the theory, with the IR fixed point (red). Arrows point towards $k \to 0$.

As we see in the right panel of Figure 3.4, the flow from all of theory space converges at one point — the *IR fixed point* of this particular theory.

Like pouring a bag of marbles down different slopes of a valley — they all end up at the bottom. Such is the picture — two completely different theories describing microscopic physics — *the UV regime* — can end up at the same *stable fixed point*[5] in the macroscopic description — *the IR regime* — which means they end up at the same *effective Hamiltonian* — their macroscopic physics behaves the same. In other words, since the partition function $\mathcal{Z}$ must contain the same information no matter the scale of the action, we get the exact same macroscopic quantities for two theories describing different microscopical physical systems, because the RG flow took them to the same fixed point. This concept is called *IR universality*.

Let us picture a slightly more involved theory, such as the one illustrated in Figure 3.5. If we start with an initial microscopic theory in the region $-1 \leq g_2 < 0,\ 0 < g_4 < 1$, we find ourselves in familiar territory — the RG flow carries us safely down to the fixed point on the left. However, stepping over the line $g_2 = 0$, the story changes: we are no longer in the *basin of attraction* of the left-hand IR fixed point. Instead, we roll down into the IR fixed point on the right. The line $g_2 = 0$ acts like a ridge in the landscape, dividing two distinct valleys — each leading to a different macroscopic theory — and destines the marble to roll down one or the other.

A much discussed concept in condensed matter physics, this is an example of a *first order phase transition* — as we cross some line of critical value, our system undergoes *spontaneous symmetry breaking* and ends up in a completely different IR fixed point,

[5] In the field of dynamical systems and chaos, this is called a *sink*.

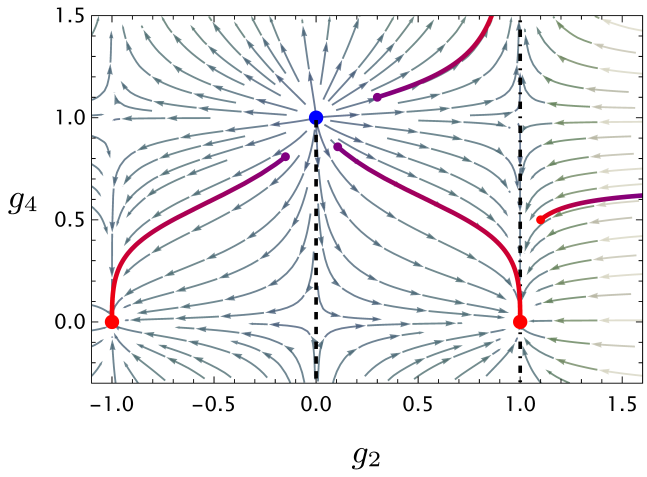

FIGURE 3.5.
RG flow in a simplified two-coupling theory space, where the arrows point towards $k \to 0$. The red dots indicate IR fixed points; the blue dot marks a UV fixed point. Trajectories flowing into the IR fixed points represent distinct low-energy theories. Any trajectory starting at $g_4 > 1$ is outside any IR fixed point basin of attraction, and will thus shoot out to infinity as $k \to 0$. The same is true for trajectories starting at $g_2 > 1$ once we let $k \to \infty$.

thereby suddenly completely changing its macroscopic properties. This is exactly the kind of mechanism which describes how water turns to ice [14].

The blue point at the top of the diagram represents a *UV fixed point*. Unlike the red IR fixed points at the bottom, where RG trajectories terminate in the IR, i.e. when $k \to 0$, the UV fixed point governs the behaviour of the theory at short distances or high energies; it is where we end up if we let $k \to \infty$, i.e. let the marble roll to a hilltop. Theories that end up at this point as we let $k \to \infty$ are said to be *asymptotically safe* — they remain predictive and well-defined all the way up to arbitrarily high energies. Crucially, only trajectories starting from a small subset of theory space can reach the UV fixed point — this subset forms the *UV critical surface*, which will be elaborated further in Section 3.2.3.

Venturing beyond the border at $g_4 = 1$, letting $k \to 0$ shoots the theory off to infinity; the RG flow no longer leads to a fixed point. The exact same behaviour is displayed if our initial theory has $g_2 > 1$ and we let $k \to \infty$ — the trajectory goes to infinity. Physically, this signals an incomplete or ill-defined description of a system. The RG flow tells us how our theory behaves at long or short distances; if it diverges off to infinity when either $k \to 0$ or $k \to \infty$, it suggests that the theory does not settle into a universal regime — we would need to step back and rethink the UV or IR description. This discussion is already suggestive of the close relationship between fixed point and renormalizability, which we will discuss in the next sections.

3.2.2 Mass Dimension, Fixed Points, and Renormalizability

To start this discussion, let us consider a generic scalar field theory with an interaction expansion:

$$\mathcal{L} = -\frac{1}{2}\phi(\partial_\mu\partial^\mu + m^2)\phi - \sum_n \frac{\tilde{g}_n}{n!}\phi^n \,, \tag{3.18}$$

where the couplings $\tilde{g}_n$ depend on the RG scale k. This is the most general form of a scalar Lagrangian, and in general, all interaction couplings are retained. From an EFT perspective, the expansion follows an ordering principle, such that the first few terms are the most important — or, *relevant* — for low-energy physics. The Wilsonian RG will allow us to understand, among other things, why EFTs work so well.

As discussed earlier, one glaring challenge in reconciling QFT with GR lies in the perturbative non-renormalizability of GR, and this issue is fuelled by the *mass dimension* of couplings. The possibility of ASQG as a solution relies on the Wilsonian idea of renormalization that we will introduce in this section.

Up until now, we have kept our treatment general and conceptual — agnostic to the number of dimensions we work in (i.e. our actions have generally been of the form $S = \int d^dx\mathcal{L}[\phi(x)]$). From now on, in view of the physical applications in the next chapters, we shall fix $d = 4$. Since we are working in natural units $\hbar = c = 1$, we define all dimensions relative to mass: $[m] = 1$, which we appropriately call the *mass dimension*.

Glancing at our Lagrangian, we see that $\partial_\mu\partial^\mu$ must have the same mass dimension as m (since all terms must have the same dimension). From this, we deduce:

$$[\partial_\mu] = 1 \,, \quad [x^\mu] = -1 \,, \quad [d^dx] = -d \,.$$

In the path integral formalism, the action must be dimensionless for the exponential to make mathematical sense, which gives:

$$[S] = 0 \,, \quad \Rightarrow \quad [\mathcal{L}] = d = 4 \,.$$

Now, the dimension of a coupling $\tilde{g}_n$ is fixed by ensuring that every term in the Lagrangian has dimension 4:

$$\tilde{g}_n \cdot m^n = m^4 \quad \Rightarrow \quad d_{\tilde{g}_n} \equiv [\tilde{g}_n] = 4 - n \,. \tag{3.19}$$

This scaling explains why GR (with Newton's constant $G_{\mathrm{N}} \sim m^{-2}$) is perturbatively non-renormalizable — it corresponds to a coupling with negative mass dimension. As discussed in Section 2.2, this causes the loops to generate divergences that require an infinite number of counterterms and, hence, an infinite number of physical observables.

But here comes a trick — we can disentangle the trivial dimensional scaling from the non-trivial running of the coupling by defining *dimensionless couplings*:

$$g_n(k) \equiv \tilde{g}_n(k)\, k^{-d_{\tilde{g}_n}} \,. \tag{3.20}$$

This rescaling strips away the dimensionality and isolates the couplings' behaviour under RG flow. It allows us to focus directly on how couplings evolve with scale, without being distracted by their dimension.

We can now rearrange the general beta function using these dimensionless couplings:

$$\beta_{g_n} = k\frac{dg_n}{dk} = \frac{dg_n}{d\log k} = (\eta[g_n] - d_{\tilde{g}_n})\, g_n(k)\,, \tag{3.21}$$

where

$$\eta[g_n] \equiv \frac{d\log \tilde{g}_n}{d\log k} = \frac{k}{\tilde{g}_n}\frac{d\tilde{g}_n}{dk} \tag{3.22}$$

is the **anomalous dimension**, capturing how quantum corrections make the system deviate from classical scaling. This is the n-dimensional version of Equation (3.17).

To derive (3.21), we note:

$$\frac{dg_n}{d\log k} = \frac{d}{d\log k}\left(\tilde{g}_n(k)\cdot k^{-d_{\tilde{g}_n}}\right) \tag{3.23}$$

$$= k^{-d_{\tilde{g}_n}}\left(\frac{d\tilde{g}_n}{d\log k} - d_{\tilde{g}_n}\tilde{g}_n\right) \tag{3.24}$$

$$= g_n\left(\eta[g_n] - d_{\tilde{g}_n}\right)\,, \tag{3.25}$$

which gives us the beta function (3.21).

We naturally expect to recover perturbation theory in the regime $|g_n| \ll 1$, i.e. when the leading order term dominates, such that we recover the classical scaling behaviour:

$$g_n(k) \approx g_n(k_0)\left(\frac{k}{k_0}\right)^{-d_{\tilde{g}_n}} \,. \tag{3.26}$$

Therefore, we expect the leading order term of $\eta[g_n]$ to be linear in g_n, such that η is negligible when $|g_n| \ll 1$. Thus, Equation (3.26) is consistent with (3.21). As discussed in Section 3.1, in QFT, the process of renormalization introduces corrections

due to integrating out high-energy modes — these are "hidden" in the anomalous dimension. From this, we see why the name "anomalous dimension" is appropriate — it quantifies the deviation from classical scaling due to quantum effects. In QFT, this would show itself as:

$$\langle \phi(x)\phi(y)\rangle \sim \frac{1}{|x-y|^{d-2+\eta}}\,, \tag{3.27}$$

instead of the classical $\langle \phi(x)\phi(y)\rangle \sim |x-y|^{-(d-2)}$, where d is the spacetime dimension of the system.

The possibility of departure from classical scaling induced by the anomalous dimension is not a minor detail, but a significant feature that leads to the development of Wilsonian renormalization. In the perturbative picture, the fate of a coupling is dictated entirely by its mass dimension, leaving only three possibilities:

- $d_{\tilde{g}n} > 0$: The coupling decreases with energy, $g_n(k) \to 0$ as $k \to \infty$; the coupling is asymptotically free.

- $d\tilde{g}n < 0$: The coupling diverges at high energies, i.e. $g_n(k) \to \infty$ as $k \to \infty$, signalling perturbative non-renormalizability, as discussed in Section 2.2.

- $d\tilde{g}_n = 0$: The coupling is classically scale-invariant (marginal), and thus remains constant across all energy scales classically.

This classification, however, is an extrapolation from perturbation theory, and is thus only valid as long as we remain in the perturbative regime $|g_n| \ll 1$. Beyond this domain, quantum fluctuations coming from loop diagrams leave their non-trivial imprint through the anomalous dimension $\eta[g_n]$. These contributions can shift the sign of the "effective scaling dimension" $\left(\eta[g_n] - d_{\tilde{g}_n}\right)$ in the beta function (3.21), changing not only the rate but crucially the sign of the exponent in a more general version of Equation (3.26). A coupling that would classically diverge may instead saturate at a finite value, or a marginal coupling may turn divergent or convergent, depending on the sign and magnitude of $\eta[g_n]$.

In this way, the anomalous dimension captures the genuinely non-perturbative effects that allow for non-trivial fixed points. Such fixed points go beyond the case of *asymptotic freedom* and provide the foundation for *asymptotic safety*: the possibility that a theory remains predictive and well-defined all the way to arbitrarily high energies. Without the anomalous dimension, the only outcomes would be trivial scaling or uncontrollable divergences. With it, new doors open to the non-perturbative regime, where new fixed points and novel scaling behaviours can emerge. To familiarise

ourselves with this regime more systematically, we now turn to a more formal analysis of fixed points and their properties.

3.2.3 Fixed Points, Stability, and Predictivity

We previously introduced the idea of RG flow and fixed points through an intuitive toy model in Section 3.2.1. Now, we turn to a more formal and quantitative analysis of fixed points and their role in defining the UV and IR structure of a theory.

In some perturbatively renormalizable theories, the coupling constants flow to zero at high energy — a property known as *asymptotic freedom* — with quantum chromodynamics (QCD) being the most well-known example of such a theory. *Asymptotic safety* generalises this idea by allowing for an interacting fixed point in the UV, where theories can flow to a fixed but non-zero value, corresponding to an interacting UV completion.

One can visualise the RG flow as motion of a marble rolling around in a mountainous landscape, where the RG scale k plays the role of elevation and the beta functions define a vector field pointing along the gradient of this potential. Dropping a marble in this space, the flow will follow the steepest descent, tracing the RG trajectories, exactly like we saw in the toy model in Section 3.2.1.

The tips and valleys of this correspond to equilibrium positions where the marble will be at a standstill — this is where the RG flow stops and it is defined by the vanishing of all beta functions:[6]

$$k\partial_k \beta_n(\{g_n^*\}) = 0\,, \tag{3.28}$$

where $\{g_n\} = (g_1, g_2, g_3, \dots)$ are the set of all couplings as, generically, beta functions are dependent on all couplings in the system. These special points are *fixed points* of the RG flow. We can identify two types of fixed points:

- **Gaussian Fixed Points (GFP)** where for all couplings $\{g_n^*\} = 0$. This corresponds to a free or non-interactive action, i.e. only the kinetic part is non-zero.

- **Non-Gaussian Fixed Point (NGFP)** where **at least one** coupling $g_n^* \neq 0$ at the fixed point.

[6] If only a subset of the beta functions vanish, we end up with a *fixed hypersurface* of dimension d_n, where d_n is the number of beta functions that does not vanish.

To identify features of our landscape, we will analyse the behaviour of the RG flow in the immediate vicinity of the fixed points. To do this, we linearise the flow around them as:

$$\vec{\beta} \simeq \vec{\beta}^* + \mathcal{J}_{\text{stab}}(\vec{g} - \vec{g}^*) \quad \rightarrow \quad \vec{g} = \vec{g}^* + \sum_i c_i \vec{e}_i \left(\frac{k}{k_0}\right)^{-\theta_i}, \tag{3.29}$$

where the stability matrix is a Jacobian matrix:

$$(\mathcal{J}_{\text{stab}})_{nm} = \left.\frac{\partial \beta_n}{\partial g_m}\right|_{g_i = g_i^*}. \tag{3.30}$$

Here $\vec{\beta}$ and $\vec{g}$ are vectors of all beta functions and couplings, θ_i are the critical exponents (minus the eigenvalues of the stability matrix), $\vec{e}_i$ are the associated eigenvectors, and c_i and k_0 are integration constants.

The eigenvalues of the stability matrix determine the nature of the fixed point:

- If $\theta_i > 0$, then $g_i(k) \to g_i^*$ as $k \to \infty$. These are **UV-attractive** or **IR-relevant** directions: small perturbations grow in the IR and thus influence long-distance physics. The blue dot in Figure 3.5 represents an example of a fixed point with two IR-relevant directions.

- If $\theta_i < 0$, then $g_i(k) \to g_i^*$ as $k \to 0$. These are **UV-repulsive** or **IR-irrelevant** directions: perturbations die away in the IR and become invisible to low-energy observers. The red dots in Figure 3.5 represent an example of fixed points with two IR-irrelevant directions.

We can, on rare occasions, have $\theta_i = 0$, in which instance $\vec{e}_i$ is called a marginal direction and then generally requires further analysis. We will, however, not deal with any marginal directions in the upcoming work. An example of RG flow displaying both a saddle GFP and an UV-attractive NGFP is displayed in Figure 3.6.

Relevant directions correspond to parameters that must be fixed experimentally. Irrelevant directions, by contrast, are fixed by the UV theory itself — their IR values are predictions rather than inputs to our theory. This division is what makes asymptotically safe theories *predictive* despite potentially containing infinitely many couplings: Provided that the number of relevant directions is finite and small, only a few couplings need to be externally fixed, while infinitely many others are predicted by the theory.

Let us build a bit of intuition for this point using our newfound language. Theory space (or the mountainous landscape) is an infinite-dimensional space spanned by all couplings, in which we, with this analysis, can identify subsets of space spanned by the relevant eigenvectors $\vec{e}_i$ with $\theta_i > 0$ of a particular UV fixed point. This subset is called

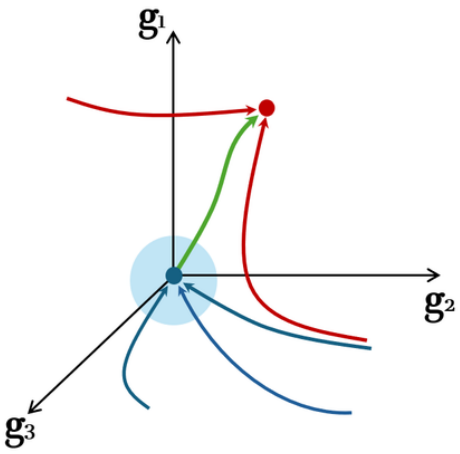

FIGURE 3.6.
Depiction of RG flow in 3-dimensional theory space of a theory with a GFP (blue dot) and a UV-completing NGFP (red dot). Arrows point toward $k \to \infty$. The GFP has two IR-relevant directions (g_2, g_3), forming a plane of asymptotically free trajectories (blue lines). The third direction (g_1) is IR-irrelevant for GFP. The NGFP also has attractive directions, giving rise to asymptotically safe trajectories (red line). The UV-complete trajectory (green line) connects the two fixed points. The blue region depicts the perturbative region, i.e. where couplings $\{g_n\} \ll 1$. Inspired by [20].

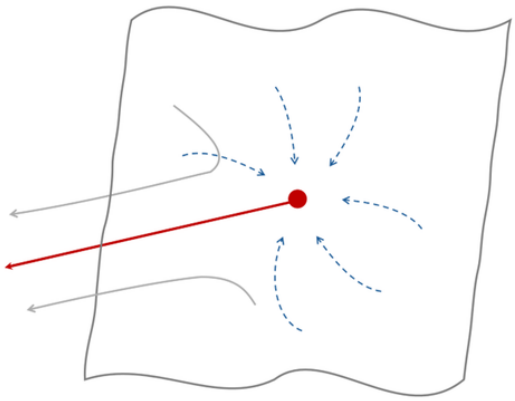

FIGURE 3.7.
The critical surface. IR-relevant/UV-attractive (blue dashed arrows) directions flow on the critical surface and approach the UV fixed point as $k \to \infty$. These directions are renormalizable. IR-irrelevant/UV-repulsive directions (grey and red arrows) leave the critical surface as $k \to \infty$, making them non-renormalizable. Inspired by [14].

the **UV critical surface** (also called the *critical manifold* and — more intuitively — the basin of attraction) and are of particular interest to us, as the set of RG trajectories that flow into a fixed point as $k \to \infty$ lives in this subspace, and thus it defines the subset of theory space for which the theory is asymptotically safe at high energies.[7] The critical manifold is depicted in Figure 3.7.

[7] We already familiarised ourselves with an IR basin of attraction in Section 3.2.1.

In other words — a theory is **asymptotically safe** if its RG trajectory remains within the UV critical surface of a fixed point — typically a NGFP — such that the couplings approach finite values as $k \to \infty$.

The number of relevant directions — i.e., the dimension of the UV critical surface — determines the number of free parameters in the theory. Thus, if this number is finite, the theory is said to be **predictive**: only a finite number of measurements are needed to determine the finite number of couplings, and thus, the RG trajectory is fully determined. The remaining infinitely many (irrelevant) directions are then simply determined by the requirement of UV safety.

In the context of quantum gravity, this is a remarkable possibility. Despite GR being non-renormalizable in the "traditional" sense of perturbative QFT, the asymptotic safety scenario suggests that the theory space spanned by the infinitely many couplings in a generic theory of gravity may contain a non-trivial fixed point with only a finite number of relevant directions. If such a fixed point exists, gravity becomes UV-complete and predictive at high energies — not by being perturbatively renormalizable, but by lying within a well-behaved RG trajectory. We will discuss the asymptotic safety scenario in detail in Chapter 4.

The Wilsonian viewpoint sets us up for practical computations using non-perturbative versions of the RG. This will be introduced in Section 3.3.

3.2.4 Wilsonian Picture and Naturalness of EFT

As we have hinted at in this chapter, the central idea behind EFT is deceptively simple: when studying physics at a particular energy scale E (equivalent to a high momentum scale), we do not need to know the full microscopic structure of the universe. If we *expect* new physics first enters above some much higher energy scale $\Lambda \gg E$, then its details are largely irrelevant to low-energy phenomena. This formulation of EFT might seem suspiciously loose. Physics is usually not the kind of science where a mere *expectation* that some higher-dimensional operator will scale as the order of 1 over some arbitrary cutoff is enough reason to justify discarding infinite towers of terms in a Lagrangian.

The Wilsonian picture of renormalization provides a perspective which makes the notion of EFTs much more natural. It considers the space of all couplings allowed by physical principles such as symmetry — see e.g. Equation (3.18) — and asks how a theory evolves as we change the scale at which it is probed [21, 22, 23]. Through formal arguments and examples, we have learned that the answer to this question

lies in the RG flow of the theory. In particular, its UV and IR behaviour are fully determined by the fixed points of the theory and their stability properties.

Starting from a UV fixed point, which defines a possible bare action for the theory, the idea of the Wilsonian approach is to integrate out momentum shells step by step, gradually removing high-energy modes and redefining the action in the process. This generates a flow in *theory space* — the space of all couplings — and under this flow, irrelevant operators are dynamically suppressed, which at the bottom rung of the scale ladder leaves us with an effective action Γ, in which the leading-order terms govern the IR dynamics. The leading-order terms of the full effective action match those that we expect from an EFT perspective.

The resulting rather simple EFT with only a few terms does not come from a hand-waving argument. It arises because the RG flow *erases* details that no longer affect observables: The low-energy effective theory is not just an approximation — it is a natural outcome of scale-dependent dynamics. EFTs are not assumptions — they are consequences of the RG.

It is, however, important to make a distinction between EFTs and effective actions. The former can only describe low-energy physics and cease to be valid beyond a certain cutoff Λ, denoting the scale of new physics. The latter can be valid at all energy scales, provided that it admits a UV completion within a QFT framework. Effective actions are thus more general and reduce to EFTs in the limit of low physical momenta.

More concretely, an EFT organises interactions in a systematic expansion of operators [16], justified by the way RG works. A generic scalar EFT could be written as

$$S_{\text{EFT}}[\phi] = \int d^4x \left[\frac{1}{2}\phi(-\partial^2 + m^2)\phi + \sum_{n=3}^{\infty} \frac{g_n}{n!} \frac{\phi^n}{\Lambda^{n-4}} \right] , \tag{3.31}$$

where we have included the renormalizable terms (operators of dimension ≤ 4 in four spacetime dimensions), and higher-dimension operators are suppressed by powers of the cutoff Λ. The couplings g_n define the *Wilson coefficients* of the theory. This expansion is analogous to a Taylor series in energy: each higher-dimension operator contributes corrections of order $(E/\Lambda)^{n-4}$, and thus becomes irrelevant at low energies. In this way, EFT can be seen as a method for "organising ignorance".

An EFT breaks down beyond the scale of new physics Λ. However, if there is no new physics, an EFT may stem from a UV completion within a QFT framework, which we may generally write down as an infinite tower of operators,

$$\Gamma[\phi] = \int d^4x \sum_n C_n \mathcal{O}_n(\phi), \tag{3.32}$$

where $\mathcal{O}$ is a set of operators to the n-th order in the given expansion scheme, and $\mathcal{C}_n$ are interaction couplings. When varying EFTs, UV-complete effective actions are valid at all energy scales, encode the momentum dependence of couplings via form factors, and the corresponding running couplings are not limited to parametrically small values. As we shall see shortly, the concept of an effective action is crucial for the formulation of non-perturbative, functional versions of the Wilsonian RG.

In the next section, we are going to formalise the concepts we learned so far, from coarse-graining to non-perturbative RG flows, beta functions, and fixed points, introducing the powerful techniques of the FRG, which make use of the concept of effective actions to describe a theory from the UV to the IR.

3.3 Functional Renormalization Group

Having introduced the Wilsonian picture of renormalization — where coarse-graining leads to scale-dependent couplings and flow trajectories in theory space — we now move to a modern framework that implements this logic at the level of the full effective dynamics: the FRG [3].

Conceptually, the FRG applies the same stepwise integration of momentum modes used in the Ising model or scalar field theory examples, but instead of tracking a few running couplings, it follows the flow of the entire effective action as quantum fluctuations are progressively included. The effective dynamics at each scale reflect the cumulative effect of integrating out modes above that scale, while lower modes are still "frozen out". The result is a scale-dependent action that evolves continuously from the microscopic theory in the UV to the full quantum effective action in the IR.

Fixed points of this flow correspond to scale-invariant regimes, just as in the Wilsonian RG, but the FRG allows us to study them beyond perturbation theory, even in strongly interacting or gravitational settings. It is this ability to capture non-perturbative behaviour, while remaining grounded in the coarse-graining logic introduced earlier, that makes the FRG a central tool in modern QFT.

To build this formalism, we will start by introducing the mathematical definition of the effective action, and then proceed to describe the shell-by-shell integration process more explicitly. We will work using the Euclidean signature, as is common in RG computations.

3.3.1 Legendre Transforms and the Effective Action

To define the effective action more precisely, we use the Legendre transform. We start from a generic bare (Euclidean) action, defined at the UV cutoff scale Λ (which we will eventually remove, by sending $\Lambda \to \infty$):

$$S_\Lambda[\phi] = \int d^4x \left(\frac{1}{2}\phi(-\partial^2 + m^2)\phi + \sum_n \frac{\tilde{g}_n}{n!}\phi^n \right) , \tag{3.33}$$

where the tilde over $\tilde{g}_n$ indicate dimensionful couplings. Restricting to field configurations with Fourier modes satisfying $|p| \leq \Lambda$, we can define the *generating functional*:

$$\mathcal{Z}[J] = \int^\Lambda \mathcal{D}\phi \, e^{-S_\Lambda[\phi] + \int J\phi} . \tag{3.34}$$

The source term J is, for our purposes, merely a practical bookkeeping tool used to generate correlation functions and does not affect the RG structure directly. This is the QFT counterpart to the partition function from statistical mechanics — it contains all information of our theory and crucially generates *all* correlation functions possible for our theory — i.e. it generates both connected and non-connected Feynman diagrams. We then define the generating functional of exclusively connected diagrams:

$$W[J] = \log \mathcal{Z}[J] , \tag{3.35}$$

and the vacuum expectation value (or "classical") field:

$$\phi_c(x) = \langle 0|\phi(x)|0\rangle = \frac{\delta W[J]}{\delta J(x)} . \tag{3.36}$$

The *effective action* $\Gamma[\phi_c]$ is then the Legendre transform of $W[J]$, where ϕ_c and J now are conjugate Legendre variables, — just like momentum and velocity for Hamiltonians and Lagrangians:

$$\Gamma[\phi_c] = \int d^4x \, J(x)\phi_c(x) - W[J] . \tag{3.37}$$

This object generates the 1PI (one-particle irreducible) correlation functions and encodes the full quantum dynamics of the theory — it "hides" all the quantum.

In the FRG framework, one modifies this definition to include a scale-dependent regulator, leading to an *effective average action* (EAA) $\Gamma_k[\phi]$ that evolves with the RG scale k and implements the feature of integrating out fast fluctuating modes shell-by-shell.

3.3.2 Shell-by-shell Integration

As we did in the treatment of the Ising model, we now split the field into low- and high-momentum modes like $\phi(x) = \bar{\phi}(x) + \delta\phi(x)$:

$$\begin{aligned}\phi(x) &= \int_{|p|\leq\Lambda} \frac{d^4p}{(2\pi)^4}\, e^{ip\cdot x}\, \tilde{\phi}(p) \\ &= \int_{|p|\leq k} \frac{d^4p}{(2\pi)^4}\, e^{ip\cdot x}\, \tilde{\phi}(p) + \int_{k<|p|\leq\Lambda} \frac{d^4p}{(2\pi)^4}\, e^{ip\cdot x}\, \tilde{\phi}(p) \\ &= \bar{\phi}(x) + \delta\phi(x)\,.\end{aligned} \tag{3.38}$$

The path integral measure correspondingly factorises as

$$\mathcal{D}\phi = \mathcal{D}\bar{\phi}\,\mathcal{D}\delta\phi\,.$$

With this decomposition, the generating functional (3.34) becomes:

$$\mathcal{Z}[J] = \int_{|p|\leq\Lambda} \mathcal{D}\phi\, e^{-S_\Lambda[\phi]+\int J\phi} \tag{3.39}$$

$$= \int_{|p|\leq k} \mathcal{D}\bar{\phi} \int_{k<|p|\leq\Lambda} \mathcal{D}\delta\phi\, e^{-S_\Lambda[\bar{\phi}+\delta\phi]+\int J(\bar{\phi}+\delta\phi)}\,. \tag{3.40}$$

To proceed, we split the bare action into:

$$S_\Lambda[\bar{\phi}+\delta\phi] = S_0[\bar{\phi}] + S_0[\delta\phi] + S_{\text{int}}[\bar{\phi},\delta\phi]\,, \tag{3.41}$$

where S_0 contains the free (quadratic) part of the action, and S_{int} contains non-trivial interaction terms, cross-terms between $\bar{\phi}$ and $\delta\phi$. This decomposition is crucial, as it isolates the part of the action that allows integration over the high-momentum fluctuations. Doing this allows us to split $\mathcal{Z}$ up into a low and high momentum part:

$$\mathcal{Z}[J] = \int_{|p|\leq k} \mathcal{D}\bar{\phi} e^{-S_0[\bar{\phi}]} \int_{k<|p|\leq\Lambda} \mathcal{D}\delta\phi e^{-S_0[\delta\phi]-S_{int}[\bar{\phi}+\delta\phi]+\int J[\bar{\phi}+\delta\phi]}\,. \tag{3.42}$$

We now focus on the second path integral in this term, which we will call a k-dependent partition function

$$\mathcal{Z}_k[J] = \int_{k<|p|\leq\Lambda} \mathcal{D}\delta\phi e^{-S_0[\delta\phi]-S_{int}[\bar{\phi}+\delta\phi]+\int J[\bar{\phi}+\delta\phi]}\,. \tag{3.43}$$

As the interactions allow for quantum fluctuations, this is where all the "quantum stuff" is kept.[8]

[8] This is what one needs to renormalize in usual perturbative renormalization.

However, performing this integration directly is highly nontrivial. Instead, we follow Polchinski's insight [18]: regulate the path integral smoothly rather than sharply cutting off momenta. This leads to the introduction of a momentum-dependent regulator $\mathcal{R}_k(p^2)$, and we define a scale-dependent deformation of the action:

$$\Delta S_k[\phi] = \frac{1}{2}\int \frac{d^4p}{(2\pi)^4}\,\phi(-p)\,\mathcal{R}_k(p^2)\,\phi(p)\,. \tag{3.44}$$

A good regulator function $\mathcal{R}_k(p^2)$ should satisfy:

- **IR suppression:** $\mathcal{R}_k(p^2) \gg p^2$ for $p^2 \ll k^2$, so low-momentum modes are effectively frozen.
- **UV transparency:** $\mathcal{R}_k(p^2) \to 0$ for $p^2 \gg k^2$, and when $k \to 0$. This ensures that high-energy modes are unaffected, and we obtain the effective action once all modes are integrated out.
- **Microscopic limit condition** $\lim_{k\to\Lambda\to\infty} \mathcal{R}_k(p^2) \to \infty$. This ensures that we regain the microscopic action in the UV.

An example of a regulator (and the one we will use in our analysis later) would be

$$\mathcal{R}_k = k^2 e^{-\frac{p^2}{k^2}}\,. \tag{3.45}$$

With the regulator in place, the k-dependent generating functional becomes:

$$\mathcal{Z}_k[J] = \int^{\Lambda} \mathcal{D}\phi\, e^{-S_\Lambda[\phi]-\Delta S_k[\phi]+\int J\phi}\,, \tag{3.46}$$

and using the Legendre transformation as in Section 3.3.1, we arrive at an expression for the *average effective action*:

$$\Gamma_k[\phi_c] = \int d^4x\, J(x)\phi_c(x) - W_k[J] - \Delta S_k[\phi_c]\,. \tag{3.47}$$

If the regulator is implemented correctly, $\mathcal{Z}_k$ in (3.46) functionally identical to $\mathcal{Z}_k$ in (3.43). As $e^{-S[\phi]-\Delta S_k[\phi]} = e^{-S[\phi]}e^{-\Delta S_k[\phi]}$ the regulator suppresses the contribution from low-momentum modes $p^2 \ll k^2$, so that the path integral effectively only includes fluctuating modes with $p^2 \gtrsim k^2$. As $k \to 0$, the regulator vanishes, all the modes get integrated out and $\Gamma_k[\phi] \to \Gamma[\phi]$, the full quantum effective action. As $k \to \Lambda$, all fluctuations are included and $\Gamma_k[\phi] \to S_\Lambda[\phi]$. This is illustrated in Figure 3.8. This construction now allows us to track how the effective action *evolves* as we gradually integrate out small-scale physics. Remember: we are not just simplifying the theory — we are tracing a *trajectory* through a space of possible theories, where the trajectory

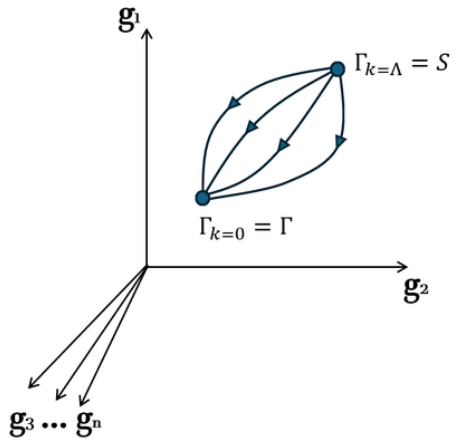

FIGURE 3.8.
FRG illustrated — both the bare action S_Λ and the effective action Γ in infinite-dimensional theory space. FRG drives $S_\Lambda \to \Gamma$ via the average effective action Γ_k. The arrows point towards $k \to 0$.

defines the *effective action* at each scale. In the next section, we introduce the FRG and the Wetterich equation, which encodes the exact RG flow of $\Gamma_k[\phi]$.

3.3.3 Wetterich Equation

We now derive the FRG flow equation for the EAA $\Gamma_k[\phi]$, first derived in [19], using a momentum-dependent regulator to suppress IR modes. We will end up with the RG flow equation in terms of the logarithmic scale derivative,

$$k\partial_k\Gamma_k[\phi] \equiv \partial_\tau\Gamma_k[\phi], \qquad \text{with} \quad \tau = \log\left(\frac{k}{\Lambda}\right) .$$

As this is not essential to the physical intuition of the project, the derivation will not be as fleshed out as other points in the thesis. We start by using the chain rule and the Legendre transformed structure of Γ_k,

$$k\partial_k\Gamma_k[\phi] = -k\partial_k W_k[J] - k\partial_k\Delta S_k[\phi] \,, \tag{3.48}$$

and from the definition of $W_k[J] = \log \mathcal{Z}_k$, we compute:

$$k\partial_k W_k[J] = -\frac{1}{\mathcal{Z}_k}\int \mathcal{D}\phi \,(k\partial_k\Delta S_k[\phi])\, e^{-S[\phi]-\Delta S_k[\phi]+\int J\phi} \tag{3.49}$$

$$= -\frac{1}{2}\int \frac{d^dp}{(2\pi)^d}\, k\partial_k R_k(p^2)\langle\phi(p)\phi(-p)\rangle_J \,. \tag{3.50}$$

Now, recalling that W_k is the generator of connected correlators,

$$\langle\phi(p)\phi(-p)\rangle_J = \frac{\delta^2 W_k[J]}{\delta J(p)\delta J(-p)} + \phi(p)\phi(-p) \,. \tag{3.51}$$

Legendre transforming and inverting the first term on the right, we get:

$$\frac{\delta^2 W_k[J]}{\delta J \delta J} = \left(\Gamma_k^{(2)}[\phi] + R_k\right)^{-1} . \tag{3.52}$$

Using this, the two-point function decomposes as

$$\langle \phi(p)\phi(-p)\rangle_J = \left(\Gamma_k^{(2)}[\phi] + R_k\right)^{-1}_{p,-p} + \phi(p)\phi(-p) .$$

Therefore, we get

$$\begin{aligned} k\partial_k W_k[J] = &-\frac{1}{2}\mathrm{Tr}\left[\left(\Gamma_k^{(2)}[\phi] + R_k\right)^{-1} k\partial_k R_k\right] \\ &+ \frac{1}{2}\int \frac{d^d p}{(2\pi)^d}\, \phi(-p) k\partial_k R_k(p^2)\phi(p) . \end{aligned} \tag{3.53}$$

Substituting back into (3.48) and noting that the disconnected part from $\langle\phi\phi\rangle_J$ cancels out $k\partial_k \Delta S_k[\phi] = \frac{1}{2}\int \phi(-p) k\partial_k R_k(p^2)\phi(p)$, we arrive at:

$$\boxed{k\partial_k \Gamma_k[\phi] = \frac{1}{2}\mathrm{Tr}\left[\left(\Gamma_k^{(2)}[\phi] + R_k\right)^{-1} k\partial_k R_k\right]} . \tag{3.54}$$

This equation governs the **exact** flow of the EAA $\Gamma_k[\phi]$ as the cutoff scale k changes, and thus, it serves as a starting point for non-perturbative analysis of QFTs using FRG.

3.3.4 Truncation Schemes

With the machinery available to us in FRG — particularly the Wetterich equation (3.54) derived in Section 3.3.3 — we in principle have the exact, non-perturbative, functional integro-differential equation for the EAA Γ_k, independent of the type of theory considered. However, as Γ_k in principle contains information on infinite terms of interaction couplings, we need some kind of approximation. Depending on the *expansion scheme* (with the vertex and derivative expansions being two examples [4]), we will have to approximate Γ_k by choosing a *truncation order*, i.e. cutting off the expansion of Γ_k to some order in the chosen expansion. We already saw an example of this in Section 3.2.1, where we truncated to four orders in the magnetisation field $m(x)$. Truncations are necessary approximations in order to solve the flow [4].

3.3.5 Closing Remarks

One of the revolutionary aspects of Wilson's formulation is that it makes no assumptions about the "true" microscopic action. There is no privileged bare action. Instead, the space of all couplings allowed by symmetry is treated democratically, and the RG flow tells us how theories evolve across scales. Think of theory space like a mountainous landscape of possible actions. In the view of the Wilsonian philosophy of renormalizaton, we do not pick a trail beforehand and hike down — we study how water flows across the whole terrain as the elevation (the RG scale k) lowers, and, crucially which mountain top the water emerges from (k going to infinity). The shape of the land (symmetries, dimensions) determines which paths are smooth, stable, or inaccessible. Thus, the Wilsonian view provides the sound underlying logic that explains what renormalization more intuitively does and why EFTs work.

4 Asymptotic Safety in Quantum Gravity

We now have a full suite of tools to describe how QFTs evolve with scale. Coarse-graining, integrating out high-energy modes, and following RG trajectories in theory space have led us to a dynamic view of QFT. As we saw in Chapter 3, renormalization is more than a computational trick — it is a conceptual paradigm shift that reframes our understanding of scale and complexity. Having laid the groundwork in the previous chapters, we are now ready to put it all to work — to climb up and down the UV/IR ladder and explore **asymptotic safety in quantum gravity**. This is the theory we will use in Chapter 6 to make predictions on black holes in Einstein-Weyl gravity, and we review it in the following.

4.1 Gravitational Effective Field Theories

Despite its UV troubles, we know that GR works beautifully as an EFT. We can see this by writing this simplified general expansion:

$$\Gamma_{\text{gravity}} = \int d^4x\sqrt{-g}\left[\frac{m_{\text{Pl}}^2}{2}R + c_1\mathcal{R}^2 + \frac{c_2}{m_{\text{Pl}}^2}\mathcal{R}^3 + \frac{c_3}{m_{\text{Pl}}^4}\mathcal{R}^4 + \cdots\right]. \tag{4.1}$$

where the first term represents the Einstein-Hilbert action known from GR, and $\{\mathcal{R}^2, \mathcal{R}^3, \mathcal{R}^4, ...\}$ represent general 2nd, 3rd, 4th (and so forth) order terms. Higher-order curvature terms are suppressed by powers of the Planck mass $m_{\text{Pl}} \sim G_{\text{N}}^{-1/2}$. These corrections first become relevant when $E \sim m_{\text{Pl}}$, which is far above any gravitational interaction we would encounter on Earth.[1]

Thus, we can safely use GR in low-energy physics such as cosmology, motion of astronomical objects and gravitational waves — because it is the leading term in an expansion.

The physical reason EFTs work so well is that quantum fluctuations at high momenta (short distances) decouple from low-energy observables — they smooth out as the magnetisation field in the Ising model. When we compute loop diagrams, the contribu-

[1] The Ricci scalar is on the order of $R \sim 10^{-33} m_{\text{Pl}}^2$ on Earth.

tions from virtual particles with momenta much greater than E manifest as corrections of higher-derivative terms suppressed by powers of $1/m_{\text{Pl}}$.

Thus, this is not just a calculational trick which makes our lives easier when calculating low-energy physics; rather, it reflects a physical structure. This insight is what justifies the use of EFTs, as discussed in Section 3.2.4.

Yet, GR ceases to work at high energies and needs to be replaced by a UV-complete theory of quantum gravity. Vice versa, starting from a given quantum gravity UV-completion, we can "flow down to IR" and see which *specific EFTs* are allowed, or, in other words, what are the bounds imposed on the Wilson coefficients by quantum gravity. In the following, we are going to focus on one such UV-completing proposal: **ASQG**.

4.2 Asymptotic Safety as a Fundamental Theory

Asymptotic safety promises a UV completion of gravity. It suggests that a gravitational theory can be UV-complete — not because it is perturbatively renormalizable, but because it flows to a *non-trivial fixed point* under the RG: it is non-perturbatively renormalizable.

There are two philosophical routes one can take when probing high-energy physics:

- A **top-down approach**, where one starts with a UV-complete theory (e.g., string theory) and flows downward to recover familiar low-energy physics. This is in principle "the ideal one" — it is long and tedious, but would (if successfully done) lead to physical predictions from first principles.[2]

- A **bottom-up approach**, where we build EFTs valid at low energies and test them. This is more of an "engineering approach". The found theories work in the regimes we would like them to, but do not promise anything about being fundamental.

Asymptotic safety — first proposed by Weinberg [1] — is a top-down approach built within the QFT framework, which naturally leads to EFT at low energies. It is based on the Wilsonian picture we discussed in the previous chapter: We need not assume

[2]From first principles: from fundamental theoretical assumptions, without relying on empirical models, phenomenological input, or approximations valid only in specific regimes.

anything about the UV or *bare* action — the RG flow encodes the information on which UV completions are possible. At the same time, it also has elements of the bottom-up approach, since its field content and spacetime dimensionality are taken to be as those we observe. In the following, we shall review the key achievements of the asymptotic safety program.

4.3 Evidence for the "Reuter Fixed Point"

The asymptotic safety scenario for quantum gravity is based on the possibility of defining gravity as a QFT, based on an interacting fixed point of the RG flow. The existence of such a UV completion for gravity — named the *Reuter fixed point* — can be systematically tested using the FRG, as outlined in the previous chapter. The idea is to start from a simple truncation and increase the truncation order step-by-step to test the stability of the results.

The simplest of such truncations is one where we essentially "stop at GR" and only include the Einstein-Hilbert term and the cosmological constant Λ:

$$\Gamma_k \simeq \frac{1}{16\pi G_{\mathrm{N}}(k)} \int d^4x \sqrt{g} \left[2\Lambda(k) - R\right] , \tag{4.2}$$

which means we have two couplings which run with the RG scale k. A milestone in the development of the asymptotic safety program was the discovery of the Reuter fixed point in the Einstein-Hilbert truncation [5]. This NGFP was obtained through the FRG analysis of the dimensionless counterparts of the couplings in (4.2) and is located at

$$g_* \approx 0.541 , \tag{4.3}$$

$$\lambda_* \approx 0.064 . \tag{4.4}$$

Linearisation around the fixed point reveals a pair of complex critical exponents,

$$\theta_{1,2} \approx 2.667 \pm 0.958i , \tag{4.5}$$

with a positive real part. This implies that the fixed point is UV-attractive and can serve as a UV completion of gravity. The existence of such a fixed point ensures that the gravitational RG flow remains finite and well-defined at all energy scales, thereby realising Weinberg's vision of a QFT of gravity based on asymptotic safety.

Expanding the truncation scheme, i.e. including more terms in the EAA Γ_k, has been of natural interest since the discovery of the Reuter fixed point. A lot of legwork

has been put into FRG analysis of more and more sophisticated truncations over the last 25 years [24, 25, 8, 26], with numerous studies indicating the existence of a non-trivial gravitational fixed point with 2 or 3 relevant directions [4]. Thus, this fixed point would give 1 or 2 free parameters, i.e. 1 or 2 observables needed to be measured for the theory to make physical predictions. More on this will be discussed in Section 4.6.1.

4.4 Connections to Standard Model

With the framework of ASQG, gravity is treated as a QFT, and therefore, testing its connections and compatibility with the SM is within reach. Intriguingly, coupling the SM to ASQG may modify the running of SM couplings near the Planck scale. This interplay opens several key avenues (see [9] for a review). For instance, gravity tends to screen gauge couplings in the UV, potentially leading to UV fixed points even for non-asymptotically free sectors (e.g., $U(1)_Y$). Secondly, the top Yukawa and the Higgs quartic couplings may be driven to fixed points by gravity, suggesting a resolution to the triviality and hierarchy problems. Thirdly, gravitational corrections to the RG flow can render previously free parameters calculable. For instance, the top and Higgs masses have been argued to emerge as a prediction [10, 27]. Finally, quantum gravity may induce a combined asymptotically safe UV completion for the SM [11].

4.5 Open Questions — Causality, Unitarity, and Existence in Lorentzian Signature

The vast majority of work done in ASQG has been in the Euclidean signature. Simply put, calculations done in the Lorentzian signature contain an imaginary unit, i, in the exponent of the path integral, rendering it oscillatory and calculations more difficult. The underlying assumption is that physical results can be recovered by Wick rotating back to Lorentzian signature at the end of the calculation, thereby saving the existence of a fixed point found in Euclidean signature [4]. Some of the work in this direction currently points to the fact that even in a Lorentzian signature, the Reuter fixed point exists and has similar properties as in the Euclidean [28].

The mere existence of a gravitational fixed point (both in Lorentzian and Euclidean signature) does not ensure ASQG to be unitary and causal. Although truncated EAAs used in practical computations sometimes show ghost poles, these are generally considered to be truncation artifacts [29]. Their residues decrease with higher-order

truncations, and they are expected to vanish in the full theory [29]. Moreover, recent literature suggests the absence of ghost modes in ASQG, once the momentum dependence of the propagator is resolved [28, 30]. Yet, more work is necessary to test unitarity and causality in asymptotic safety. In particular, as we shall discuss later, positivity bounds need to be checked.

4.6 Climbing Down the Ladder — Landscapes in Quantum Gravity

The existence of the fixed point is a cornerstone in ASQG, but what does this mean for the physics far below the UV regime? This is a natural question to ask, and the answer is almost "built in" to the machinery of asymptotic safety, where we let $k \to 0$ and end up with an effective action. Even more picturesque, we end up in the *asymptotic safety landscape*.

The terminology of "landscape" and "swampland" originally emerged in the context of string theory [31], where the vast number of low-energy EFTs seemingly compatible with quantum gravity are separated into those that can be consistently UV-completed ("landscape") and those that cannot ("swampland"). In recent years, this language has been adopted in the context of ASQG [20, 32, 33] to describe the space of effective actions generated by RG flows from a UV fixed point, as depicted in Figure 4.1. The idea is conceptually similar: only certain combinations of Wilson coefficients correspond to consistent, UV-complete theories. This motivates introducing the notion of an "asymptotic safety landscape", which we now describe.

4.6.1 Landscapes in Asymptotic Safety

As discussed in Section 4.1, the effective action encodes an infinite tower of curvature terms, each accompanied by a Wilson coefficient. Within a given UV completion, some Wilson coefficients will be fixed by the underlying theory. For instance, in ASQG, the Wilson coefficients are determined via the RG flow emerging from the NGFP in the UV. The space of all effective actions that are reachable via asymptotically safe RG flows defines the *asymptotic safety landscape*. Each point in this landscape corresponds to an EFT. In practical calculations, each EFT in the landscape is associated with the $k \to 0$ limit of an FRG trajectory starting from the fixed point, with the corresponding Wilson coefficients given by the IR values of the running couplings $G_i(k)$ [20, 32].

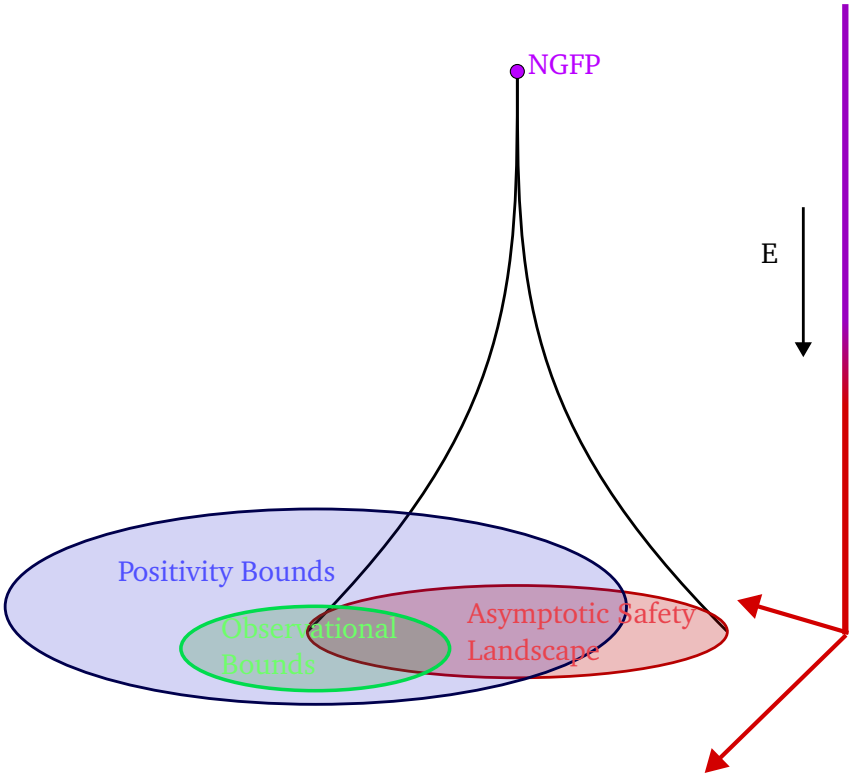

FIGURE 4.1.
Sketch of the idea behind the concept of asymptotic safety landscape. The asymptotic safety landscape (red set) emerges from UV-complete trajectories ending up at the NGFP, thus mapping out the viable EFTs of ASQG. Bounds such as positivity bounds (blue set) or observational bounds (green set) may further constrain the landscape of physically viable EFTs.

If a fixed point has N relevant directions, it gives rise to an $(N-1)$-dimensional sub-landscape from that fixed point — the dimension is reduced by one because the RG scale can always be absorbed by fixing a unit of mass (e.g., G_{N}, as we will do in Chapter 6). The total asymptotic safety landscape is then the union of these sub-landscapes across all viable UV fixed points. This also indicates the landscape emerging from the gravitational fixed point discussed in Section 4.3 would have 1 or 2 free parameters, therefore needing 1 or 2 observables to be measured for the theory to be completely determined.

In this thesis, we will see a concrete realisation of this idea. As we shall see in Chapter 6, the NGFP of Einstein-Weyl gravity has 1 relevant direction, and is therefore a 0-parameter theory. The Wilson coefficient of the Weyl-squared term, G_{C^2}, thus becomes a derived quantity once asymptotic safety is imposed.

These concepts connect to our discussion in Section 3.2.3 and Section 4.1. Just as higher-order operators in gravity are suppressed by powers of the cutoff scale, the EFTs within the asymptotic safety landscape inherit constraints from their UV origin. Only those combinations of Wilson coefficients that come from an asymptotically safe

RG trajectory (or from another UV completion of gravity) are physically allowed — all others lie in the swampland of inconsistent EFTs [31].

4.6.2 Intersection with Positivity Bounds and Additional Consistency Constraints

Beyond the requirement of UV completion, not every EFT in the landscape is automatically physically viable and additional consistency constraints must be checked. One prominent example is *positivity bounds*, which originate from the fundamental principles of unitarity, locality, analyticity and Lorentz invariance [34, 35]. These bounds constrain the signs and magnitudes of Wilson coefficients in the effective action and can rule out large classes of EFTs that otherwise appear consistent.

Other criteria, such as causality, analyticity of correlation functions and compatibility with black hole thermodynamics [36], can also impose nontrivial restrictions on the effective action. These bounds further constrain the asymptotic safety landscape (and perhaps other quantum gravity landscapes) on top of the positivity bounds [32]. Finally, observational bounds may also constrain quantum gravity landscapes. These ideas are illustrated in Fig. 4.1.

4.6.3 From Landscapes to Quantum Spacetime

The power of the asymptotic safety landscape program [20, 32] lies not only in the possibility of identifying consistent EFTs, but also in providing a direct bridge between fundamental microphysics and the macroscopic structure of quantum spacetime. By computing the RG flow from first-principle calculations, one can find which quantum-corrected spacetimes are physically realisable within a UV-complete theory of gravity. A central goal of this thesis is to demonstrate the power of this idea by illustrating how the requirement of UV completion can constrain, and in some cases — like the results of this thesis — uniquely determine, the form of gravitational solutions in the IR.

5 Classical Einstein-Weyl Gravity Phase diagram

In this chapter, we will discuss higher-derivative corrections to Einstein gravity and how they impact black hole physics. Specifically, we shall review the classical "phase diagram" of Einstein-Weyl gravity [15], which maps out the set of possible static, spherically symmetric black hole solutions which arise when quadratic corrections are added to the Einstein-Hilbert truncation. Importantly, this analysis is purely classical and does not treat the action as a QFT; we are temporarily setting aside quantum effects to understand the space of classical solutions beyond GR. This knowledge will then be used in combination with the asymptotic safety analysis in the next chapter to determine which black hole-like solutions can stem from an asymptotically safe UV completion in the Einstein-Weyl approximation.

5.1 Interlude: Quadratic Gravity

Before diving into the specific solutions of Einstein-Weyl gravity and the resulting phase diagram, it is worth taking a step back to motivate our setup. As discussed in Section 2.2, the Einstein-Hilbert action on its own is famously perturbatively non-renormalizable when treated as a QFT, due to the need for infinitely many counterterms, making the theory unpredictive in the UV. One natural way to try to fix this is by expanding the action to include higher-order terms; indeed, historically, this was one of the first attempts to consistently quantize gravity, going beyond quantum GR [37]. The first such corrections are the second-order curvature terms, leading to a theory known as quadratic gravity. The general action reads:

$$S = \int d^4x\sqrt{-g}\left[\frac{1}{16\pi G_{\mathrm{N}}}R + \alpha R^2 + \beta R_{\mu\nu}R^{\mu\nu}\right], \tag{5.1}$$

where we have ignored the topological Gauss–Bonnet term, which does not contribute to the equations of motion in four dimensions. This action can be rewritten in terms of the Weyl tensor as:

$$S = \int d^4x\sqrt{-g}\left[\frac{1}{16\pi G_{\mathrm{N}}}R + \frac{1}{2\lambda}C^2 - \frac{\omega}{3\lambda}R^2\right], \tag{5.2}$$

with the square of the Weyl tensor $C^2 = C_{\mu\nu\rho\sigma}C^{\mu\nu\rho\sigma}$ capturing the traceless part of the Riemann tensor. This action represents the most general, local, diffeomorphism-invariant modification to GR up to second order in derivatives.

In 1977, Stelle showed that this theory is not just better-behaved than GR — it is perturbatively renormalizable [37]. The reason is simple: the higher-derivative terms improve the UV behaviour of the propagator, which was the driving force behind the superficial degree of divergence being $D > 0$. Instead of falling off like $1/p^2$, it now falls off like $1/p^4$, which tames the divergences that plague quantum GR. With this, quadratic gravity became the first known QFT of gravity that satisfies power-counting renormalizability. However, *this gain of perturbative renormalizability comes at a serious cost*.

When linearised around flat space, the theory propagates with more degrees of freedom than GR. In addition to the usual massless graviton of Einstein-Hilbert gravity, the spectrum includes:

- A massive spin-0 scalar (coming from the R^2 term),
- A massive spin-2 mode (from the $R_{\mu\nu}^2$ or Weyl-squared term).

The scalar mode is not necessarily problematic. The issue lies with the massive spin-2 field. Its kinetic term carries the wrong sign, making it a *ghost* — a particle with negative norm. This induces a violation of unitarity and leads to instabilities, as the vacuum could decay into ghost–normal particle pairs, and the vacuum energy is not bounded from below anymore.

Summarising, on one hand, quadratic gravity is a rare example of a quantum gravity theory that is well-behaved in the UV. On the other hand, it is haunted by a fundamental problem — lack of unitarity — which makes the theory physically unviable.

Despite this, the quadratic extension of GR is the only one so far where it was possible to **map out the set of possible black hole-like solutions at a classical level**. From a fundamental physics perspective, such a quadratic extension may be regarded as a ***truncation* of the full effective action** stemming from a unitary and UV-complete theory — which would remove the ghost from the spectrum, e.g., via a non-trivial momentum dependence of the propagator [29]. This is, for instance, the case in ASQG [28, 30], and this is the philosophy we are using in our work.

5.2 Vacuum Solutions and the Gravitational Phase Space

In the following, we review the results of [15], which constructed the classical phase diagram of static, spherically symmetric vacuum solutions in Einstein–Weyl gravity.

We work within the quadratic truncation, using the philosophy underlined in the previous section and setting the R^2 and cosmological constant terms to zero. This reduces the action to:

$$\Gamma = \frac{1}{16\pi G_{\rm N}} \int \mathrm{d}^4x\sqrt{-g}\left(R - \frac{1}{2}G_{\rm C^2}\, C^{\mu\nu\rho\sigma}C_{\mu\nu\rho\sigma}\right), \tag{5.3}$$

where $C^{\mu\nu\rho\sigma}$ is the Weyl tensor. Note that we are using the notation Γ instead of S because in the next chapter we shall interpret it as a truncation of the full effective action; yet, in [15] the action above is used as a classical action. Varying this action with respect to the inverse metric as

$$\frac{\delta\Gamma}{\delta g^{\mu\nu}}$$

yields the equations of motion, which we now set out to analyse following the procedure of [15].

5.2.1 Long Distance Solutions

Just like if we were to derive the Schwarzschild solution of the Einstein-Hilbert action (i.e. the Weyl-squared term turned off), we consider vacuum solutions of Einstein-Weyl gravity with a static, spherically symmetric geometry described by the metric:

$$ds^2 = -h(r)dt^2 + \frac{dr^2}{f(r)} + r^2 d\Omega^2\,, \tag{5.4}$$

where the functions $h(r)$ and $f(r)$ are determined by solving the equations of motion arising from the Einstein-Weyl action. As we look for asymptotically flat solutions at large distances, we express the metric far away from a black hole solution using the weak field limit:

$$h(r) = 1 + \epsilon V(r), \qquad f(r) = 1 + \epsilon W(r)\,, \tag{5.5}$$

which leads to the solutions at large distances:

$$\begin{aligned} h(r) &\sim 1 - \frac{2G_{\mathrm{N}}M}{r} + 2S_2^- \frac{e^{-m_2 r}}{r}\,, \\ f(r) &\sim 1 - \frac{2G_{\mathrm{N}}M}{r} + S_2^- \frac{e^{-m_2 r}}{r}(1 + m_2 r)\,, \end{aligned} \tag{5.6}$$

where M is the ADM mass, $m_2 = 1/\sqrt{G_{\mathrm{C}^2}}$ is the mass of the spin-2 mode introduced by the Weyl-squared term, and S_{-2} is an integration constant — the so-called Yukawa charge. In this solution, 2 terms were omitted. One term is $\sim S_2^+ e^{m_2 r}$, and goes away due to the assumption of asymptotic flatness, while the other one gets absorbed into the definition of the ADM mass [38, 39]. Note that G_{C^2} is now referred to as some mass, rather than a coupling/Wilson coefficient. It may seem obvious — if we are abandoning the notions of QFT, it seems natural not to denote it as a coupling. However, something deeper is going on. A glaring weakness of quadratic gravity is the massive "ghost" particles the quadratic terms bring with them. We are not immediately worried, as quadratic gravity in the eyes of an effective field theorist or an asymptotic safety practitioner is merely a truncation of higher-order gravity. Thus, the ghost particles are expected to be artefacts of this specific truncation scheme and will cease to exist in a full expansion of gravity, where the Wilson coefficient still will be retained [4, 29].

The gravitational field is thus characterized by M and S_{-2}, once G_{N} and m_2 are fixed. These parameters define a two-dimensional phase space, in which different classes of solutions occupy distinct regions. This yields the *phase diagram* of Einstein-Weyl gravity. But for now, m_2 is still a free parameter.

5.2.2 Close-Up Solutions and Frobenius Analysis

To formulate finite radius solutions, one can expand the metric functions near finite or vanishing radius using a generalised Frobenius method [38, 15] — the generalisation of a Taylor expansion for non-regular expressions (which black hole solutions famously are):

$$\begin{aligned} h(r) &= (r - r_0)^t \left[\sum_{n=0}^{N} h_{t+\frac{n}{\Delta}} (r - r_0)^{\frac{n}{\Delta}} + \mathcal{O}\left((r - r_0)^{\frac{N+1}{\Delta}} \right) \right], \\ f(r) &= (r - r_0)^s \left[\sum_{n=0}^{N} f_{s+\frac{n}{\Delta}} (r - r_0)^{\frac{n}{\Delta}} + \mathcal{O}\left((r - r_0)^{\frac{N+1}{\Delta}} \right) \right]. \end{aligned} \tag{5.7}$$

Solutions are classified by their exponents $(s,t)^{\Delta}_{r_0}$, and a complete catalogue is available in Table I in [15]. For instance:

- $(1,1)^1_{r_H}$ corresponds to black holes,

- $(-1,-1)^1_0$ to repulsive naked singularities,
- $(-2,2)^1_0$ to attractive naked singularities,
- $(1,0)^2_{r_T}$ to non-symmetric wormholes.

A version of this was first done in 2015 [38]. The Frobenius expansion near $r = 0$ (centre) or $r = r_H$, r_T (horizon/throat radius) of a black hole-like solution provides not only the classification of the different families, but also a starting point for numerical integration.

5.2.3 The Shooting Method

To bridge the large-distance expansion (5.6) with the Frobenius-based close-up expansion, a numerical shooting method has been employed in the literature [40, 41, 42, 38, 39, 15]. Starting from the weak-field expansion at large r, and using the Frobenius series as an inner boundary condition near $r = 0$ or the throat/horizon radius r_0, a boundary value problem emerges. The shooting method uses numerical integration to stitch together the ends of the solutions, thereby mapping the landscape of input parameters of the large-distance solutions of the metric (M and S_{-2}) to the corresponding families of black hole-like solutions, i.e. **black hole mimickers**. Further details on this are found in [15].

5.3 Structure of the Phase Diagram

The space of solutions can be organised into the discussed distinct families [15], each corresponding to a qualitatively different gravitational structure:

- **Black holes:** Given by the $(1,1)^1_{r_H}$ expansion. These include both Schwarzschild and non-Schwarzschild black holes and form a one-dimensional manifold in the phase space. Turning off the Weyl-squared term $G_{\mathrm{C}^2} = m_2^{-2} \to 0$ turns off the Yukawa term $S_2^- e^{-m_2 r} \to 0$, which takes us to the Schwarzschild solutions and lands us on the dashed line of the phase diagram.
- **Type I (Naked Singularities):** Characterized by divergent $h(r)$ near $r = 0$, these solutions are approximated by the $(-1,-1)^1_0$ Frobenius family. They dominate large parts of the phase diagram and have curvature invariants which scale as $R_{\mu\nu}R^{\mu\nu}$ and $R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma}$ diverging as $\sim r^{-6}$. This is also the solution of negative-mass black holes in classic GR, which can intuitively be seen by “going

left" from the Schwarzschild black hole solutions (dashed line) into the negative ADM mass regime in the phase diagram, where we land in the Type I region. It attracts ($\partial_r h(r) < 0$) in the regime $r > 2M$ like a standard black hole, but the divergent nature near $r = 0$ shifts the sign ($\partial_r h(r) > 0$) such that it acts repulsively when $r < 2M$.

- **Type II (Bachian Singularities):** These vanish at the origin and belong to the $(-2, 2)^1_0$ family. They feature strong attractive potentials and curvature invariants diverging like r^{-8}. Like the Type I solution, they are horizonless solutions with singularities, with the added feature that the Bach tensor diverges as we tend to $r \to 0$, and are therefore only meaningful objects in quadratic gravity.[1]

- **Geometric asymmetry of Type III wormholes:** By far the most exotic of the black hole mimickers, these solutions belong to the $(1, 0)^2_{r_T}$ Frobenius family and resemble wormholes, which are connecting two regions of spacetime. Like normal wormholes, the throat is defined by $f(r_T) = 0$, but $0 < |h(r_T)| < \infty$, implying the absence of a horizon *and* the possibility of traversing it and appearing on the other side. However, unlike standard wormholes, they are not symmetric across the throat. One side of the geometry is asymptotically flat, while the other asymptotes to a rapidly decaying (and going to a singular) geometry. In standard traversable wormholes, the geometry is symmetric across the throat: Both sides are asymptotically flat as the metric functions smoothly approach Minkowski space as $r \to \pm\infty$ [44]. By contrast, Type III wormhole solutions in Einstein-Weyl gravity are asymmetric: One side of the throat is asymptotically flat for $r \gg r_T$, mimicking Schwarzschild or Minkowski behaviour, while the other side does *not* approach flat space. Instead the metric functions decay exponentially as $r \to \infty$:

$$h(r) \sim e^{-r}\,, \qquad \frac{1}{f(r)} \sim e^{-r}$$

and "funnels" the geometry into a compressed region of spacetime, which leads to a singularity. The lack of symmetry implies a directional nature: signals and geodesics can pass through from one side but experience drastically different conditions depending on which side they enter. This asymmetry means the wormhole connects one familiar, extended universe to a highly curved or even singular region, unlike symmetric wormholes, which connect two mirror asymptotic regions.

[1] The Bach tensor is a traceless fourth order tensor:$\mathcal{B}_{\mu\nu} \equiv \left(\nabla^\alpha\nabla^\beta + \frac{1}{2}R^{\alpha\beta}\right) C_{\mu\alpha\nu\beta}$. In conformal gravity, the Bach tensor plays the role the Einstein tensor does in GR [43].

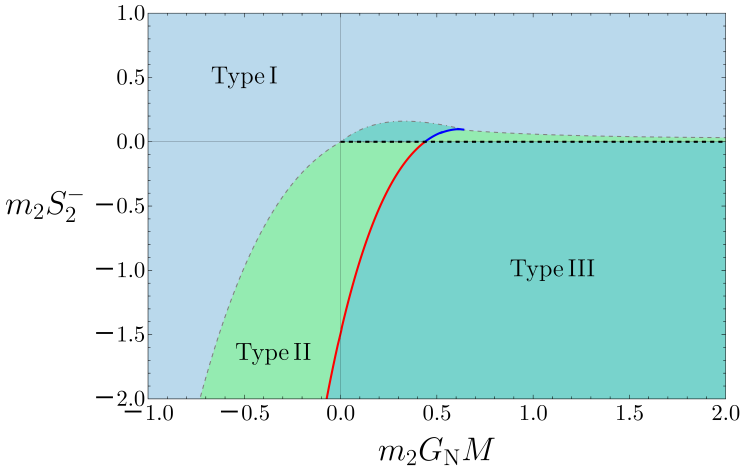

FIGURE 5.1.
Classical phase diagram of Einstin-Weyl gravity, where m_2 is a free parameter, as seen in [15]. The families of solutions are mapped as a function of the ADM mass M and Yukawa charge S_2^-. Type I: Naked Singularities. Type II: Bachian naked singularities. Type III: asymmetric Wormholes. Red and Blue lines: Non-Schwarzschild black holes. Dashed line: Schwarzschild black holes. Flat Minkowski space is at the origin.

Figure 5.1 [15] shows these families mapped according to the parameters (M, S_{-2}), with "real" black hole solution[2] lines and triple points marking transitions between families, analogous to phase diagrams/flow diagrams in condensed matter physics. The Minkowski vacuum sits at the origin, while the *massive triple point* occurs at $M \approx 0.623$, $S_2^- \approx 0.102$, indicating a secondary *non-Minkowski vacuum*.

The phase diagram also suggests that most generic vacuum configurations ($M = 0$) correspond to naked singularities or wormholes as the S_2^--axis spans these solution families almost exclusively. It is important to stress that this is not a "real" phase diagram where, for instance, triple points indicate a scale-independent theory with second-order transition and critical behaviour — it is merely a great (and fitting for this thesis) analogy.

In the next chapter, we will, after painstakingly deriving the Weyl-squared Wilson coefficient — which determines m_2 — feed this quantum gravity determined value into this diagram, identifying how asymptotic safety constrains the set of possible black hole-like solutions in the Einstein-Weyl approximation.

[2] "Real" as in having a horizon.

6 Black hole Mimickers from Asymptotically Safe Einstein-Weyl Gravity

After traversing the ideas behind ASQG, we now apply them to a concrete setting — the Einstein-Weyl truncation — with the ultimate goal of providing a proof of principle for the construction of the landscape of quantum black holes stemming from a given UV completion of gravity.

6.1 Idea: Black Holes from Quantum Gravity

As we have seen, the simple addition of the Weyl-squared term massively enriches the space of black hole solutions with much more exotic objects. This is true in general: higher-derivative corrections add on more solutions to the field equations; the Einstein-Weyl approximation is the simplest correction to GR and the only one where the classical phase diagram of solutions has been completely traced out [15]. Contrary to what the 2-dimensional depiction of the phase diagram might communicate at first glance, it needs 3 input parameters to identify the allowed gravitational object — M, S_2^- and m_2, the latter "pointing out of the page" in Figure 5.1, and essentially setting the scale of the diagram. This thesis is based on the cornerstone idea that quantum gravity, and in particular ASQG, is expected to predict some of the Wilson coefficients in the low-energy EFT. Since $m_2 = 1/\sqrt{G_{C^2}}$[1], we can now use the machinery of the RG to compute the asymptotic safety landscape to fix the value of G_{C^2}, and hence m_2. Indeed, as we shall see, the Einstein–Weyl truncation admits an asymptotically safe UV completion, and the corresponding landscape is a 0-parameter theory. What was previously treated as a free parameter is thus determined from first principles, allowing us to identify the correct slice of phase space selected by asymptotic safety and constrain the space of possible solutions compatible with a UV-complete gravitational theory.

[1] This is defined as $m_2^2 = \frac{\gamma}{2\alpha}$ in the original paper [15].

6.2 Fixed Points of Einstein-Weyl Gravity

The analysis in this section is based on the β-functions derived in [8]. We will use these beta functions to derive the flow in the Einstein-Weyl subspace and to extract a prediction for the corresponding Wilson coefficient.

6.2.1 Dimensionless Couplings and Beta Functions

Einstein-Weyl gravity can be viewed as a special case of higher-derivative gravity. The beta functions for the full quadratic theory have been computed in detail in [8]. The Einstein-Weyl truncation is a 2-dimensional subset of the 4-dimensional theory space of full quadratic gravity, in which only the Weyl-squared coupling G_{C^2} is kept nonzero:

$$\Gamma = \frac{1}{16\pi G_N}\int \mathrm{d}^4x\sqrt{g}\left(R - \frac{1}{2}G_{C^2}\, C^{\mu\nu\rho\sigma}C_{\mu\nu\rho\sigma}\right) .$$

Contrary to Chapter 5, the sign in front of the determinant of the metric g is now positive, as the FRG analysis is carried out in an Euclidean signature.

In four dimensions, both the Newton coupling G_N and the Weyl-squared coupling G_{C^2} have mass dimension $[G_N] = [G_{C^2}] = -2$. To study their flows, these couplings must be rendered dimensionless by rescaling them with the RG scale k with the mass dimension $[k] = -1$:

$$g(k) := G_N(k)\, k^2 , \qquad g_{C^2}(k) := G_{C^2}(k)\, k^2 . \tag{6.1}$$

We also introduce the logarithmic RG "time":

$$\tau(k) := \log\left(\frac{k}{k_0}\right) \tag{6.2}$$

for later convenience. The quantity k_0 is an integration constant setting the scale of reference of our flow, $\tau(k_0) = 0$. The RG flow is then described by two beta functions:

$$\beta_g[g, g_{C^2}] := \frac{dg}{d\tau} , \quad \beta_{C^2}[g, g_{C^2}] := \frac{dg_{C^2}}{d\tau} . \tag{6.3}$$

As a first step of our analysis, we need to understand whether these beta functions possess any fixed points that could serve as a UV completion for the truncated theory.

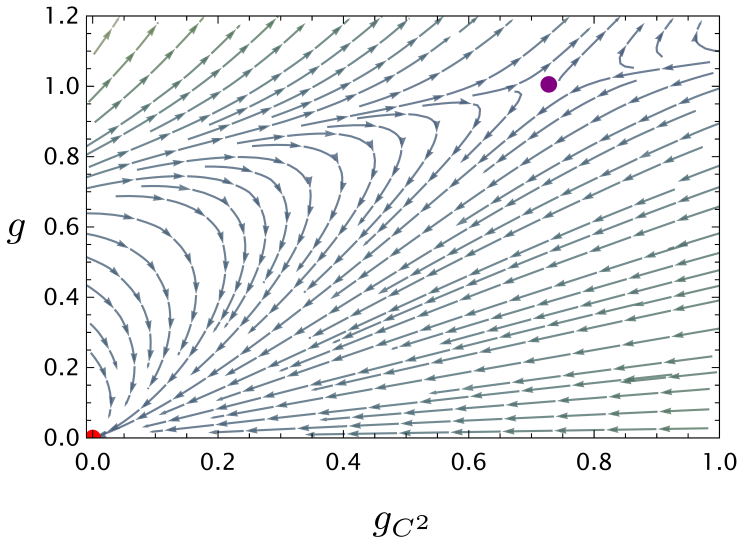

FIGURE 6.1.
Flow of Einstein-Weyl gravity with one GFP (red) and one interacting NGFP (purple). The flow of the arrows points towards decreasing RG scale k.

6.2.2 Flow Diagram and Fixed Point Analysis

As we discussed in Chapter 3, the beta functions predict how the theory flows through theory space, like how a marble rolls around in a mountainous landscape. The fixed points — i.e. where the marble will be at a standstill — are defined by the vanishing of both beta functions:

$$\beta_g[g^*, g^*_{\mathrm{C}^2}] = 0\,, \quad \beta_{\mathrm{C}2}[g^*, g^*_{\mathrm{C}^2}] = 0\,. \tag{6.4}$$

Solving these equations yields two fixed points:

- A **GFP** at $\{g^*, g^*_{\mathrm{C}^2}\} = \{0, 0\}$, corresponding to a free theory in the deep IR.

- A **NGFP** at $\{g^*, g^*_{\mathrm{C}^2}\} = \{1.0053, 0.7277\}$, corresponding to an interacting theory in the UV.

These fixed points shape the dynamics of the RG flow, as illustrated in the RG flow diagram in Figure 6.1. Glancing at this flow diagram, it seems entirely possible to reach the GFP if one just drops a marble immediately below the NGFP.

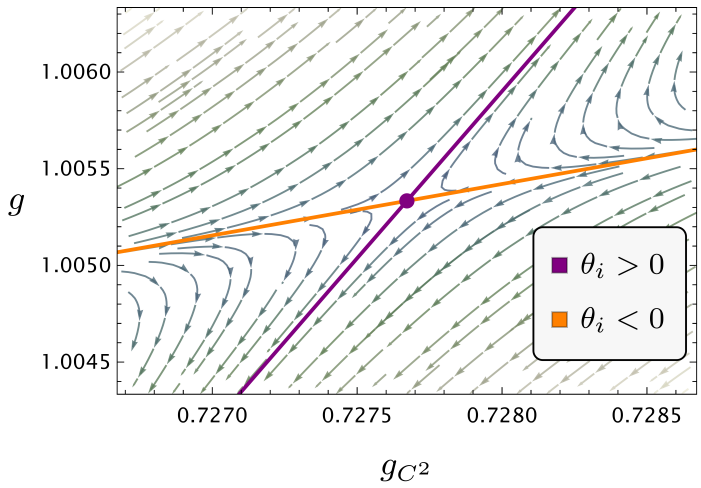

FIGURE 6.2.
Flow near the NGFP with eigendirections overlaid. The purple line denotes the IR-relevant direction; the orange line denotes the IR-irrelevant direction. The arrows point towards $k \to 0$. Note: The eigenvectors (6.5) are stacked as $\vec{e}_i = \begin{bmatrix} g \\ g_{C^2} \end{bmatrix}$, why the axes may appear flipped.

To characterise the behaviour near the NGFP, we linearise the flow around it and diagonalise the Jacobian stability matrix to compute the critical exponents and corresponding eigenvectors:

$$\theta_1 = 2.617\,, \quad \vec{e}_1 = \begin{bmatrix} -0.865 \\ -0.501 \end{bmatrix}\,, \qquad \theta_2 = -0.937\,, \quad \vec{e}_2 = \begin{bmatrix} -0.257 \\ -0.966 \end{bmatrix}\,. \tag{6.5}$$

Referring to Equation (3.29), these results show that the NGFP has:

- One IR-relevant direction ($\theta_1 > 0$) — perturbations grow as $k \to 0$ and affect low-energy physics.

- One IR-irrelevant direction ($\theta_2 < 0$) — perturbations shrink as $k \to 0$ and become unobservable at low energies.

In other words, any trajectory flowing out of the NGFP along the IR-relevant direction corresponds to a UV-complete theory — but only one such trajectory exists. This behaviour is evident if we zoom in on the NGFP, as visualised in Figure 6.2, where the eigendirections have been plotted along with the RG flow. When placing an initial condition exactly on the IR-irrelevant (orange) direction, the RG trajectories flow towards the NGFP as $k \to 0$. The opposite is true for the IR-relevant (purple) eigendirection: the NGFP is reached as $k \to \infty$.

The fixed point data are summarised in Table 6.1. This shows that the NGFP provides a unique direction for constructing a UV-complete theory — one which flows toward

	g^*	$g^*_{C^2}$	θ_i
GFP	0	0	$\{-2, -2\}$
NGFP	1.0053	0.7277	$\{-0.937, 2.617\}$

TABLE 6.1.
Table summarising the fixed point analysis.

the IR along the relevant direction while suppressing deviations along the irrelevant one.

6.2.3 Finding the Ultraviolet Complete Trajectory

As previously speculated, it is quite easy to make the marble go to the IR valley of the GFP — just place it below the orange line, and it will roll safely to the IR. However, we are not just searching for some EFT which works reasonably well, we are looking for one coming from a UV-complete theory — this presents a far greater challenge. As the NGFP has only one relevant direction, we need to push the ball up the mountain in **exactly** the right direction and make sure it ends at the very top of it, hitting the NGFP. Landing at the top, we have found *asymptotic safety*.

The *unique* RG trajectory that flows from the NGFP in the UV down to the GFP in the IR is the so-called *separatrix*. Since there is only one IR-relevant direction, the NGFP defines a 0-parameter theory:[2] Once the trajectory is fixed, the corresponding RG running is uniquely determined. Solving the beta functions as a non-linearly coupled system of differential equations gives us the exact flow of the theory from the far UV ($\tau \to \infty$) to the IR ($\tau \to -\infty$).[3] The challenge is to find the right initial conditions of the flow corresponding to the only asymptotically safe trajectory of our system.

By perturbing around the NGFP in theory space ever so slightly along the IR-relevant direction, we get an initial coupling point $\{g(\tau_0), g_{C^2}(\tau_0)\}$, from which we numerically integrate "upwards" (letting $\tau \to \infty$). If we choose the wrong direction, the trajectory diverges — "rolling off into infinity". Then we can take steps around the NGFP and reiterate the process of integrating upwards until we hit the NGFP. Once we choose the correct initial condition, we land (up to numerical precision) on the one true separatrix and connect the NGFP to the IR. This will be our strategy.

This procedure is depicted in Figure 6.3. The computed separatrix remains stable over a range $\tau - \tau_0 = 100$, corresponding to more than 43 orders of magnitude in scale: $k/k_0 > e^{100} \sim 10^{43}$. This provides strong numerical evidence that it is the

[2]The number of parameters of a theory is $N - 1$ due to one setting the scale of the system [4], as explained in Section 4.6.1.
[3]Note that $\tau \to \infty$ is equivalent to $k \to \infty$, and $\tau \to -\infty$ is equivalent to $k \to 0$.

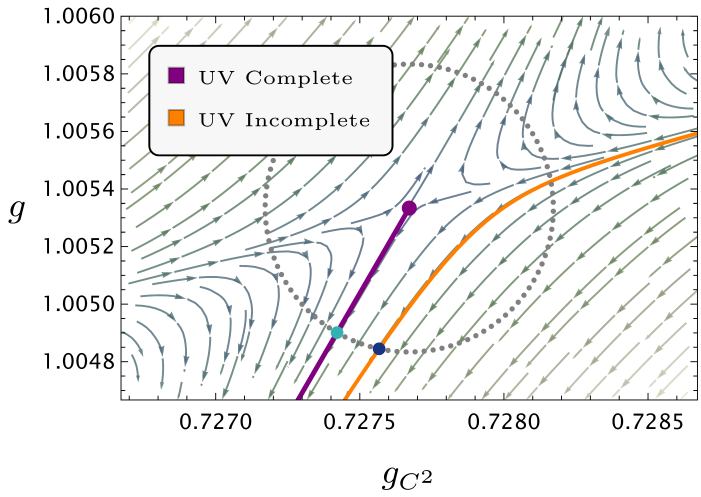

FIGURE 6.3.
Illustration of the method of finding the unique UV-complete RG trajectory. Starting from a slightly wrong initial condition (dark blue dot), the RG trajectory misses the NGFP and shoots off to infinity. Starting on the right trajectory (light blue dot), the RG flow connects to the NGFP in the UV.

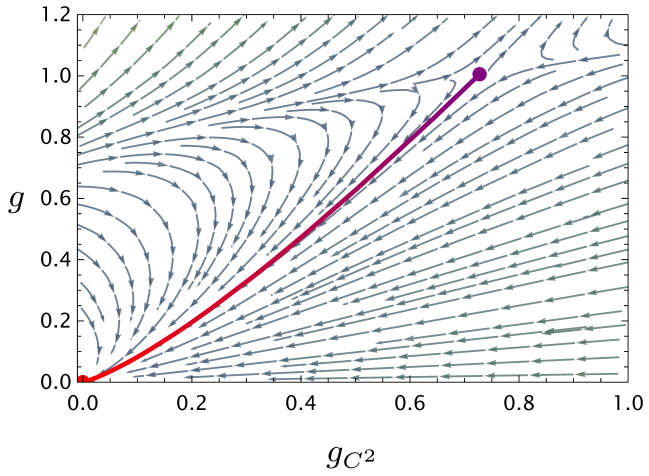

FIGURE 6.4.
RG flow of the theory overlaid by the unique UV-completing trajectory connecting the NGFP and the GFP.

"right" UV-complete trajectory. We thus end up with the only UV asymptotically safe trajectory as shown in Figure 6.4. Since the UV-complete RG trajectory is unique, it yields a 0-parameter theory: once we are on the right separatrix, there is no freedom left — the RG flow will carry us up and down the k scale ladder and safely land us in the UV or IR in a completely determined fashion. In the next section, we will use this UV-complete trajectory to extract the Wilson coefficient of the Weyl-squared term.

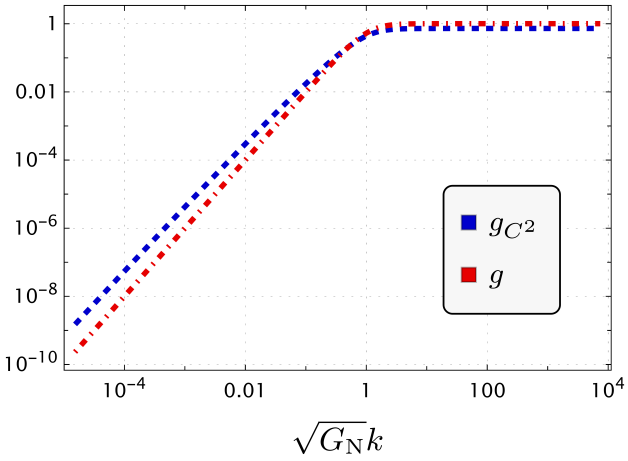

FIGURE 6.5.
Plot of the two couplings in the Einstein-Weyl truncation. For $k \to \infty$, they approach their NGFP value. As $k \to 0$, they scale according to their mass dimension + some logarithmic running for $g_{C^2,\mathrm{IR}}$. As both have the same mass dimension, the logarithmic running explains why they do not run parallel to each other as $k \to 0$.

6.3 Determining the Wilson Coefficient from Asymptotic Safety

As we make our way down the RG scale ladder, we go toward the IR or EFT regime — all quantum fluctuations have been filtered out, and we are left with an EFT description of our theory determined by the Wilson coefficients — i.e., how it would look like for someone living on earth in a low gravity/coupling regime.

6.3.1 Infrared Limit and Logarithmic Running

As $k \to 0$, the dimensionful couplings tend to their Wilson coefficient. By that logic, we would expect the dimensionless couplings to scale as their mass dimension $[G_{\mathrm{N}}(k)] = [G_{\mathrm{C}^2}(k)] = -2$, like

$$\{G_{\mathrm{N}}(k), G_{\mathrm{C}^2}(k)\} \to \text{constant} \qquad \Rightarrow \qquad g(k) \propto k^2\,, \quad g_{\mathrm{C}^2}(k) \propto k^2\,. \tag{6.6}$$

As both couplings have the same mass dimension, we would expect them to be parallel in a logarithmic plot in the IR regime $k \to 0$. However, looking at Figure 6.5, we see that this is not the case — almost, but not quite.

As it turns out, the culprit causing this non-parallel behaviour is graviton fluctuations. As described in [20, 32], the presence of massless fluctuations (i.e. the graviton) gen-

erates a logarithmic IR running in $g_{C^2}(k)$. We should therefore update our expectation of the running near $k = 0$ to fit the generic ansatz:

$$g^{(\mathrm{IR})}(k) \simeq g_{\mathrm{IR}} \left(\frac{k}{k_0}\right)^2 , \quad g^{(\mathrm{IR})}_{C^2} \simeq g_{C^2,\mathrm{IR}} \left(\frac{k}{k_0}\right)^2 \left[1 - b \log\left(\frac{k}{k_0}\right)\right] , \tag{6.7}$$

where $\{g_{\mathrm{IR}}, g_{C^2,\mathrm{IR}}, b\}$ are constants. This is leaving us in a bit of a pickle; due to the properties of logarithms and the arbitrary nature of k_0,

$$\log\left(\frac{k}{k_0}\right) = \log\left(\frac{k}{\tilde{k}_0}\frac{\tilde{k}_0}{k_0}\right) = \log\left(\frac{k}{\tilde{k}_0}\right) + \log\left(\frac{\tilde{k}_0}{k_0}\right) . \tag{6.8}$$

This means that we can always come up with some redefinition $\tilde{k}_0 = k_0 \cdot k_1$ which keeps us from filtering out the logarithmic running and finding the constant IR-value of $g_{C^2,\mathrm{IR}}$. If we make such redefinitions, the IR-value $g_{C^2,\mathrm{IR}}$ changes while keeping the term $g_{C^2,\mathrm{IR}} \cdot b$ constant — our 0-parameter theory has, in the blink of an eye, seemed to turn into a 1-parameter theory with an arbitrary, unphysical scale as our parameter. However, using some physical value from our real world to hold onto, we can get rid of this pesky new parameter and *separate the wheat from the chaff*.

As we approach the IR, we know that the Newton coupling must approach Newton's constant G_{N}:

$$G_{\mathrm{N}} = \lim_{k\to 0} G_{\mathrm{N}}(k) = \lim_{k\to 0} g^{(\mathrm{IR})}(k) k^{-2} = g_{\mathrm{IR}} k_0^{-2} . \tag{6.9}$$

This can be used as an input parameter to fix the scale k_0 such that

$$g^{(\mathrm{IR})}(k) = g_{\mathrm{IR}} \left(\frac{k}{k_0}\right)^2 = G_{\mathrm{N}} k^2 \qquad \Leftrightarrow \qquad k_0^2 = G_{\mathrm{N}}^{-1}, \text{ where } g_{\mathrm{IR}} = 1 . \tag{6.10}$$

Doing this, we have chosen G_{N} to fix the arbitrary scale/integration constant k_0, which in terms of gravitational RG flow has the interpretation of the transition scale to the quantum gravity regime. Since k_0 is arbitrary but the same in the whole system, we can rewrite $g_{C^2}(k)$ in the IR as

$$\begin{aligned} & g^{(\mathrm{IR})}_{C^2}(k) = g_{C^2,\mathrm{IR}} \left(\frac{k}{k_0}\right)^2 \left[1 - b \log\left(\frac{k}{k_0}\right)\right] , \\ \xrightarrow{\text{fix } k_0} \quad & g^{(\mathrm{IR})}_{C^2}(k) = g_{C^2,\mathrm{IR}} G_{\mathrm{N}} k^2 \left[1 - b \log\left(\sqrt{G_{\mathrm{N}}} k\right)\right] . \end{aligned} \tag{6.11}$$

Using our scale-fixed IR-expressions for the couplings, we can now consider the ratio

$$\frac{g^{(\mathrm{IR})}_{C^2}(k)}{g^{(\mathrm{IR})}(k)} = g_{C^2,\mathrm{IR}} \left[1 - b \log\left(\sqrt{G_{\mathrm{N}}} k\right)\right] = g_{C^2,\mathrm{IR}} [1 - b\tau] , \tag{6.12}$$

where now we have chosen k_0 such that $\tau(1/\sqrt{G_{\mathrm{N}}}) = 0$, namely the RG time vanishes at the Planck mass, since $G_{\mathrm{N}} = m_{\mathrm{Pl}}^{-2}$ in natural units. Once the scale k_0 is specified, we

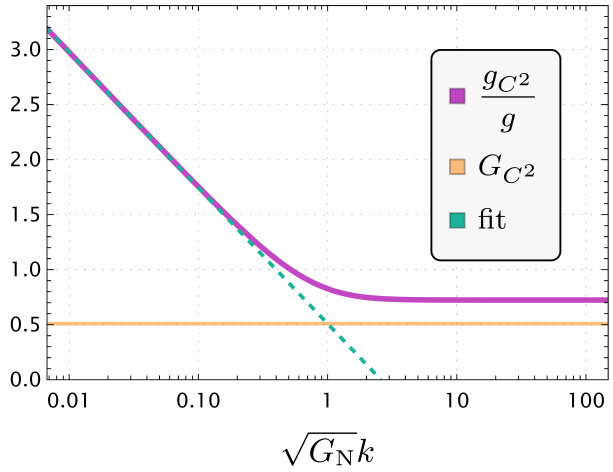

FIGURE 6.6.
Fit which determines the Wilson coefficient. After fixing the parameter $k_0 = \frac{1}{\sqrt{G_N}}$, the slope of $g^{(IR)}_{C^2}(k)/g^{(IR)}(k)$ determines the Wilson coefficient of the Weyl-squared term G_{C^2} at the intersect $\sqrt{G_N}k = 1 \Leftrightarrow \tau = 0$.

have a well-defined prescription to single out the value of $g_{C^2,IR}$, just by subtracting the logarithm present in Equation (6.12). This means that the only remaining task is to determine the slope b. To do this, we use the fit function in Mathematica to fit to the linear (linear in τ) function (6.12)

$$g_{C^2,IR} = \lim_{k\to 0}\left[\frac{g_{C^2}(k)}{g(k)} - k\frac{\partial}{\partial k}\left(\frac{g_{C^2}(k)}{g(k)}\right)\log\left(\sqrt{G_N}k\right)\right], \tag{6.13}$$

where $k\frac{\partial}{\partial k}\left(\frac{g_{C^2}(k)}{g(k)}\right) = -b$. Now the Wilson coefficient of the Weyl-squared coupling in the Einstein-Weyl action is given simply by:

$$G_{C^2} = g_{C^2,IR}\, G_N\,. \tag{6.14}$$

Doing the calculation in Mathematica, the Wilson coefficient given in Planck units is found:

$$g_{C^2,IR} = 0.5092 \quad\Longrightarrow\quad \boxed{\mathbf{G_{C^2} = 0.5092\ m_{Pl}^{-2}}}\,. \tag{6.15}$$

This process is demonstrated in Figure 6.6.

6.3.2 Interpreting the Fit and Fixing the Scale

It is easy to get lost in the mathematical details when subtracting the logarithm, so let us create some intuition for this process. A helpful way to visualise the determination of the Wilson coefficient is through the ratio $\frac{g^{(IR)}_{C^2}}{g^{(IR)}(k)}$. From Equation (6.12), we see that the ratio is an affine function which depends linearly on the RG time τ, and its intercept at $\tau = 0$ (i.e. $\frac{k}{k_0} = e^{\tau=0} = 1$ on the logarithmic plot in Figure 6.6) gives

us $g_{C^2,\mathrm{IR}}$. However, there is a subtlety: the location of the axis $\tau = 0$ is arbitrary due to the arbitrariness of k_0. This means that we can shift the entire plot along the horizontal axis, depending on how we define the quantum gravity scale k_0 — or, in other words, scale the values on the horizontal axis.

By fixing the RG time to vanish at the Planck scale, $\tau(1/\sqrt{G_\mathrm{N}}) = 0$, we effectively centre the plot such that $\frac{k}{k_0} = 1$ when $k = 1/\sqrt{G_\mathrm{N}}$. In this frame, the fit intersects the vertical axis precisely at the physical Wilson coefficient we seek. In physical terms, this means identifying the quantum gravity scale k_0 with the Planck mass.

On a more practical level, when numerically solving the coupled system of differential equations in Mathematica, we use the FRG time τ and specify initial conditions $\{g(\tau_0), g_{C^2}(\tau_0)\} = \{g_0, g_{C^2,0}\}$ (where we have made sure that we pick a point from the UV-complete trajectory). Changing the value of $k_0 \to \tilde{k}_0 = \alpha k_0$ (with α being an arbitrary dimensionless number) is equivalent to shifting the RG time origin:

$$\tau_0 \to \tau_0 + c\,, \quad \text{with} \quad \{g(\tau_0 + c), g_{C^2}(\tau_0 + c)\} = \{g_0, g_{C^2,0}\}\,. \tag{6.16}$$

This shift modifies the value of $g_{C^2,\mathrm{IR}}$, but preserves the physical content of the theory:

$$\tilde{g}_{C^2,\mathrm{IR}} \cdot \tilde{k}_0^{-2} = (\alpha^2 g_{C^2,\mathrm{IR}}) \cdot (\alpha k_0)^{-2} = g_{C^2,\mathrm{IR}} \cdot k_0^{-2} = g_{C^2,\mathrm{IR}} \cdot G_\mathrm{N} = G_{C^2}\,. \tag{6.17}$$

As expected, $\sqrt{g_{C^2,\mathrm{IR}}}$ must vary inversely with k_0 to maintain this relation. This has been checked numerically and holds across all choices of k_0, confirming that the extracted Wilson coefficient is well-defined. This will also become clearer in the following section, where we shall add more details on the numerical analysis performed.

6.4 Interlude: Numerical Improvement

As already mentioned, the system we have worked with, whose beta functions are provided in [8], is solved numerically. This comes with a set of challenges and pitfalls which one should be aware of. In this section, we will go into several of the factors which have been considered and improved, focusing on the details that matter for our scope.

6.4.1 How Mathematica Handles Numbers

Mathematica works with different types of numbers, which broadly can be put into two categories:

- **Exact numbers** — numbers such as intergers, fractions, irrational numbers like π, $\sqrt{2}$ and e are all examples of this. These are known down to infinite decimals and used in symbolic calculations, where they pose no rounding errors

- **Floats** — these are inexact and only known down to some decimal point as a number of bits (usually 16 decimal places which requires 53 bits[4]) are used to store the information. As an example, if we ask Mathematica to store a number with 3 decimal places, $0.834 \neq 0.834000$ — Mathematica does not know about the information after the 3 first decimals.

This is something we should take into account when we make numerical calculations in Mathematica. When we solve the system of differential equations, Mathematica is (according to some method it chooses — e.g., Euler's method) numerically integrating once for every minuscule step in the trajectory. This is iterated many times over, which maps out a trajectory like the one in Figure 6.4. As numerical methods utilise floats, Mathematica needs to allow some error for every calculation, i.e. every small step in the trajectory.

This motivates us to discuss accuracy and precision in Mathematica:

- **Precision** — The number of correctly significant digits — e.g., 2.65, 0.00538 and $402 \cdot 10^9$ are all numbers with precision of 3.

- **Accuracy** — The number of correct digits after the decimal point — e.g. 1543.794 and 0.001 are both numbers with accuracy of 3.

For every calculation/step in the trajectory, Mathematica will try to achieve a number within some error tolerance called *AccuracyGoal* and *PrecisionGoal*. It stops and goes onto the next calculation/step when either goal is satisfied, meaning the least stringent of the two conditions determines the actual error tolerance [45]. Imagine working with AccuracyGoal and PrecisionGoal both set to 6. When working with numbers ~ 1, this is no problem, as the result of each calculation is within 0.001%. But going down to numbers $\lll 1$ (e.g., result $= 2 \cdot 10^{-6}$), the result fulfils the accuracy requirement; thus, Mathematica stops and goes on to the next calculation. This error suddenly

[4] As bits work in binary, 16 decimals require $10^{16} = 2^{n_b} \rightarrow n_b = \log_2(10^{16}) \approx 53.15$ number of bits

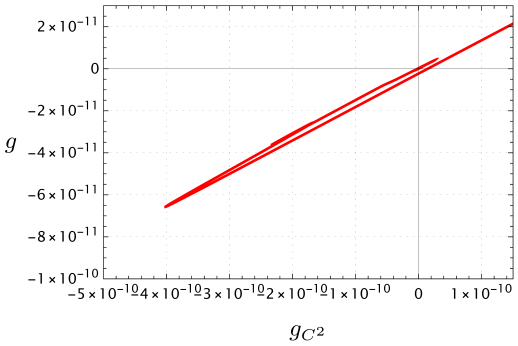

(A) Normal accuracy: Well-behaved down to 10^{-10}.

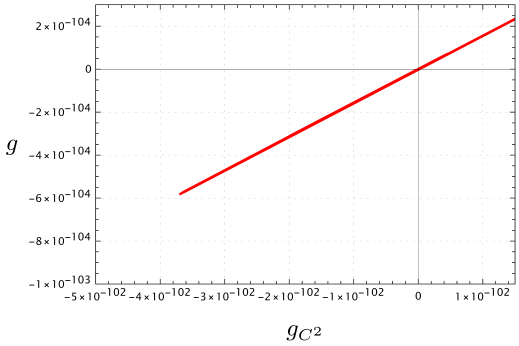

(B) Improved accuracy: Well-behaved down to 10^{-103}.

FIGURE 6.7.
Comparison of trajectories before and after improving accuracy. Around the origin, which is the region where we have extracted the Wilson coefficient G_{C^2}, the accuracy has been improved by over 90 orders of magnitude.

makes a substantial difference proportionally, as the uncertainty of the calculation is now within the same order of magnitude as the actual result. As our work concerns the IR, we are in the regime around these small numbers $\lll 1$.

6.4.2 Increased Accuracy in the Infrared Regime

To mitigate these problems, AccuracyGoal and PrecisionGoal were adjusted. We can see the significance of these adjustments in the *before and after* plots of the trajectory around the GFP/origin of theory space in Figure 6.7.

From the IR behaviour in expression (6.7) we know that this trajectory should not be able to cross any of the axes and venture into negative couplings — the beta functions should simply not allow it. Thus, any trajectory which wanders off into negative couplings can only be a result of numerical inaccuracy within the error tolerance. Surely enough, increasing AccuracyGoal lowers the magnitude of these "origin fluctuations" as we would like — over 90 orders of magnitude in the Figure 6.7.

These origin fluctuations cause mathematical singularities in the ratio (6.12) (due to the coupling $g(k)$ in the denominator approaching and crossing 0), which are essential for the calculation of the slope b via a fitting function that allows us to filter out the logarithmic running and thereby determine G_{C^2}. The effect of this increase in accuracy can be seen in Figure 6.8.

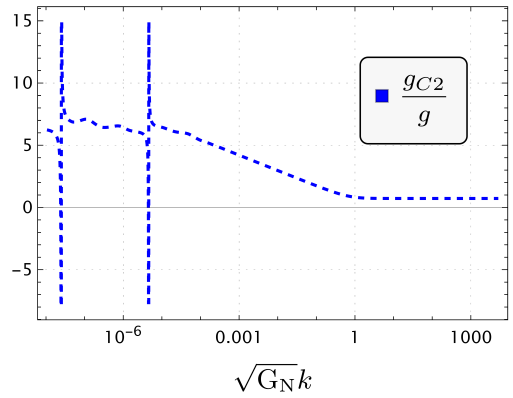

(A) Normal accuracy: The lack of accuracy around the origin creates numerical "singularities" when extracting the Wilson coefficient.

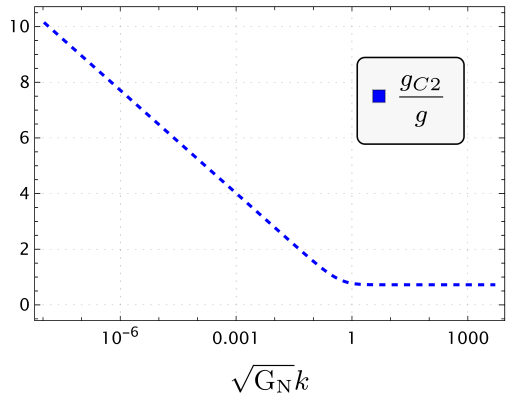

(B) Improved accuracy: The ratio now behaves smoothly.

FIGURE 6.8.
Comparison of the ratio g_{C2}/g. After improving the accuracy, there are no more numerical "singularities" — therefore, we can trust the function fitted to the slope and the value of G_{C^2}.

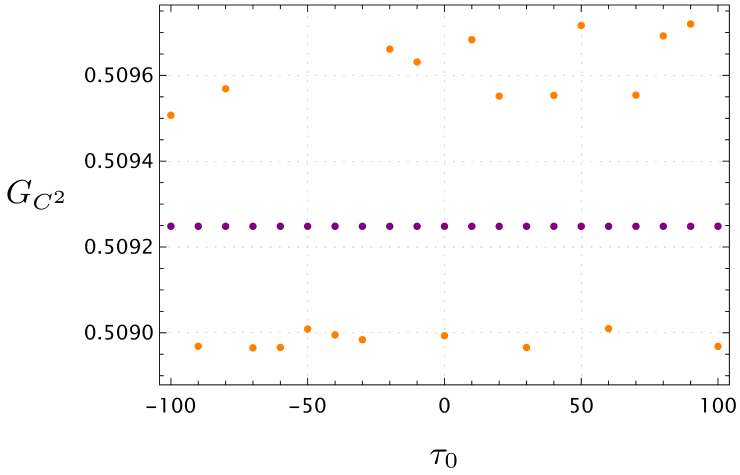

FIGURE 6.9.
Scatterplot of obtained values for the Wilson coefficient G_{C^2} as a function of the initial FRG time τ_0. The found values are centred around a constant value of G_{C^2}, confirming a consistent subtraction of the logarithmic running. The orange dots represent unaltered accuracy in Mathematica, whereas the purple ones represent improved accuracy.

6.4.3 Accuracy of the Result

As discussed in Section 6.3.2, we can use a sanity check on the fit function to confirm that we have found the correct Wilson coefficient of this truncation.

Doing the calculations for a slew of different $\tau_0 = [-100, 100]$, our sanity check confirms the method, as we get the same result (within some uncertainty), as depicted by the orange dots in Figure 6.9. Improving the accuracy, we see how the result converges to be (almost) on one line (purple dots) — the uncertainty has been decreased significantly.

Summarising the calculations, the analysis and numerical improvements outlined in this section yield the Wilson coefficient:

$$G_{\mathrm{C}^2}^{macc} = 0.5093 \pm 3.2 \cdot 10^{-4}\, m_{\mathrm{Pl}}\,, \xrightarrow{\text{Improve accuracy}} G_{\mathrm{C}^2}^{acc} = 0.509248283 \pm 1.8 \cdot 10^{-9}\, m_{\mathrm{Pl}}\,. \tag{6.18}$$

The numerical uncertainty could have been improved even further, but this machinery is a proof of concept — the exact number down to, e.g., 30 decimal places is not necessarily more interesting to us.

Having discussed numerical improvement and found the prediction on the Wilson coefficient G_{C^2}, we are now ready to see how this result impacts the classical analysis in [15]. Before that, an interlude on the validity of the analysis in [15] is in order.

6.5 Interlude: Domain of Validity of the Classical Phase Diagram

In addition to our ASQG computation, which sets the value of the Weyl-squared Wilson coefficient G_{C^2}, we re-examine the derivations in [15], identifying additional constraints on its validity. These constraints will appear in the ASQG-induced slices of the 3D phase diagram in the form of blurred-out regions, as we will see in Section 6.6.

It is a striking fact that adding just a single term to the Einstein–Hilbert action massively enriches the space of solutions with exotic objects as naked singularities and asymmetric wormholes. But to make real sense of what this enrichment tells us — and which parts of it might actually correspond to physical solutions — we need to constrain the allowed gravitational parameters and carefully select the regions of validity of the analysis in [15].

The weak-field solution (5.6) used in [15] to derive the classical phase diagram (cf. Chapter 5) assumes that all contributions to the metric are small and of comparable order:

$$\frac{G_{\mathrm{N}} M}{r} \sim \frac{S_2^- e^{-m_2 r}}{r} \sim S_2^- m_2 e^{-m_2 r} \ll 1\,. \tag{6.19}$$

These assumptions imply two facts:

- The expansion is valid for $r \sim 1/m_2$
- $G_{\mathrm{N}} M \sim S_2^- \ll 1/m_2$

However, since the Schwarzschild metric is an exact solution of the full equations, we could argue that this constraint can be relaxed when studying configurations near the Schwarzschild (dashed) line of $M \geq 0,\ S_2^- = 0$. This slightly extends the region of validity, which can then be represented pictorially as a blurred-out region, rather than a sharp constraint.

A more refined approach would be to expand around the Schwarzschild background for Yukawa perturbations:

$$\begin{aligned} h(r) &\sim 1 - \frac{2G_{\mathrm{N}}M}{r} + \epsilon V(r)\,, \\ f(r) &\sim 1 - \frac{2G_{\mathrm{N}}M}{r} + \epsilon W(r)\,, \end{aligned} \tag{6.20}$$

This would enable a consistent treatment of solutions close to Schwarzschild:

$$h(r) \sim 1 - \frac{2G_{\mathrm{N}}M}{r} + \epsilon \sum_n V^{(n)}(r)\,, \qquad f(r) \sim 1 - \frac{2G_{\mathrm{N}}M}{r} + \epsilon \sum_n W^{(n)}(r)\,. \tag{6.21}$$

Performing a Frobenius type analysis (like in Chapter 5) on V and W, we get a solution which reduces to (5.6) after a redefinition of the ADM mass M with one subtle difference: now the M does not come from a linear expansion, and is therefore not subject to a constraint — but we still impose that the (5.6) are valid for $r \sim 1/m_2$ and $S_2^- \ll 1/m_2$.

This formalism ensures that Schwarzschild solutions remain within the valid domain, and the magnitude of Yukawa corrections controls the deviations. The key takeaway is that the phase diagram is most reliable near the Schwarzschild region, and therefore, we do not impose a cutoff on the phase diagrams as we go along the M-axis, rather only as we deviate orthogonally from it. If one deviates significantly, higher-order corrections become essential to the analysis, and the confidence in the mapping of the families of solutions decreases.

6.6 Einstein-Weyl Phase Diagram from Asymptotic Safety

By applying the framework of ASQG to Einstein-Weyl gravity, we have extracted a Wilson coefficient directly from first-principle calculations. This gives us the ability to overlay our quantum gravity input onto the classically derived phase diagram, thus modifying it and constraining the solution space.

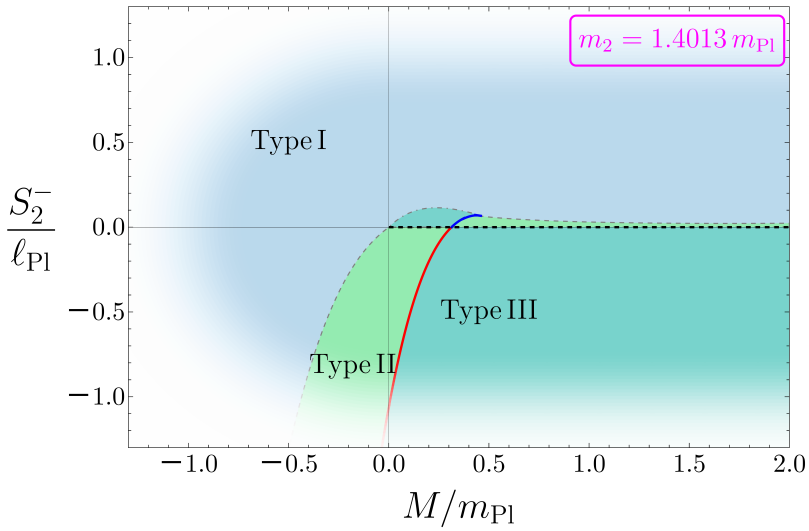

FIGURE 6.10.
Modified (M, S_2^-) phase diagram with the asymptotic safety value $m_2 = 1.4013\, m_{\rm Pl}$. This is a rescaled version of the diagram discussed in Chapter 5, now scaled to a quantum gravity-determined value. Type I (naked singularities) and Type III (non-symmetric wormholes) dominate. Type II (Bachian singularities) occupy a suppressed region.

Previously, in [15] and in the previous chapter, the parameter m_2 had been fixed to some convenient value in Planck units. But in reality, the original diagram is just a 2D slice of a fully three-dimensional phase space (M, S_2^-, m_2), where the m_2 would be the axis "pointing out of the paper". With our quantum gravity-derived value of $G_{\rm C^2}$, we can now replace that arbitrary choice with a physically selected value:

$$m_2 = \frac{1}{\sqrt{G_{\rm C^2}}} = \sqrt{1.9637}\, m_{\rm Pl} = 1.4013\, m_{\rm Pl}\,. \tag{6.22}$$

Importantly, in constructing the new phase diagrams, we will also account for the new constraints we derived in Section 6.5. These, as anticipated, will translate into blurred-out regions in the ASQG-induced phase diagram.

We begin by inserting the above value of m_2 into the original (M, S_2^-) diagram. Since m_2 essentially sets the scale, the resulting ASQG diagram, seen in Figure 6.10, is a zoomed-out version of Figure 5.1. This is more properly interpreted as the asymptotic safety-determined slice through the 3D phase space. To understand the impact more clearly, we examine slices of the full diagram from different orientations. Figure 6.11 shows the phase diagram projected into the (m_2, S_2^-) plane — i.e. turned on its side — at fixed ADM mass $M = m_{\rm Pl}$. Here, we can clearly see the slice our asymptotic safety analysis has predicted, depicted by the vertical magenta line representing a fixed m_2. This clearly intersects the Type I and Type III regions, which heavily dominate, while almost completely avoiding Type II altogether.

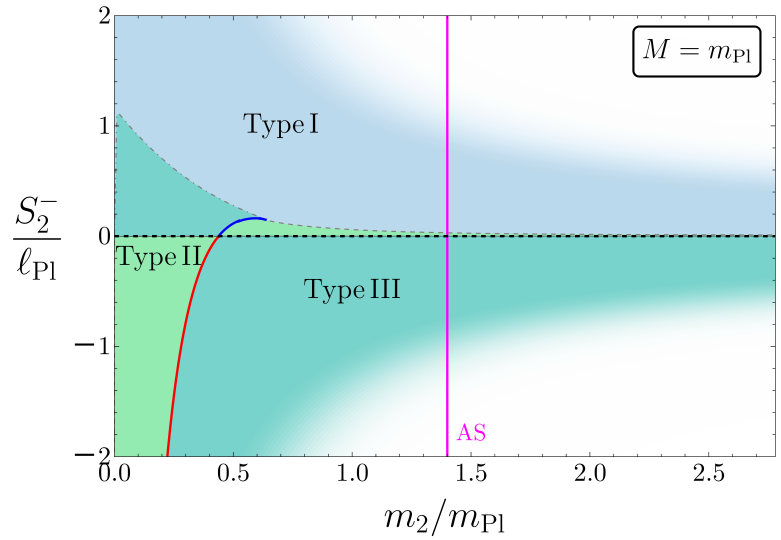

FIGURE 6.11.
Phase diagram slice in the (m_2, S_2^-) plane at $M = m_{\rm Pl}$. The magenta line marks the asymptotically safe value $m_2 = 1.4013\, m_{\rm Pl}$, which intersects the Type I and III regions and disfavours the Type II branch.

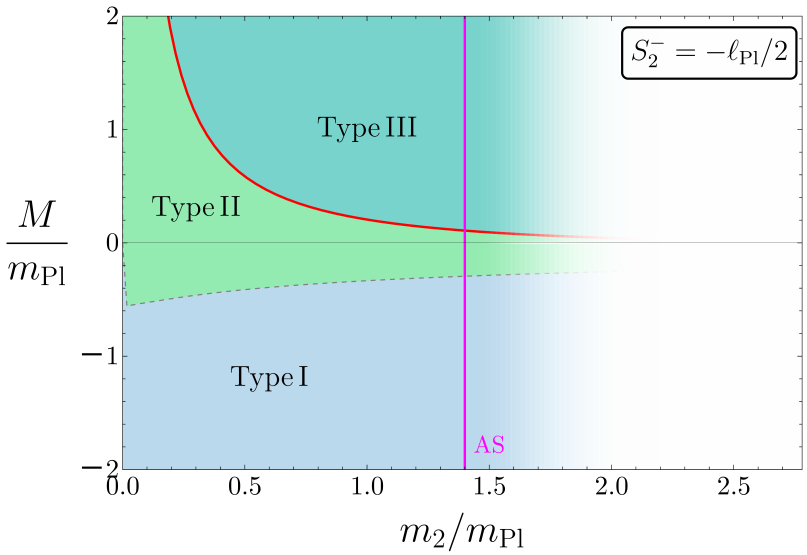

FIGURE 6.12.
Phase diagram slice in the (m_2, M) plane at fixed $S_2^- = -\ell_{\rm Pl}/2$. The asymptotic safety line intersects Type I and III regions, suggesting that these solutions are consistent with the UV-completion of Einstein-Weyl gravity.

Looking at the original diagram from the bottom (and rotated 90 degrees in a counter-clockwise position), we examine the (m_2, M) plane. Figure 6.12 fixes the Yukawa charge at $S_2^- = -\ell_{\rm Pl}/2$. Once again, Type I and III regions dominate in the asymptotic safety slice, although not as heavily as in Figure 6.11. Finally, Figure 6.13 shows the same (m_2, M) plane at $S_2^- = \ell_{\rm Pl}/2$. In this case, the asymptotic safety value lands in a region where only type I is a valid family of solutions.

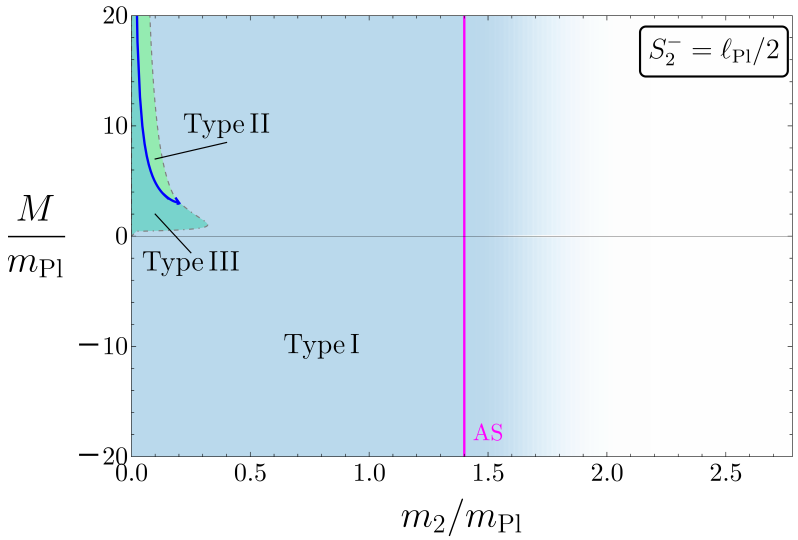

FIGURE 6.13.
Phase diagram slice at $S_2^- = +\ell_{\rm Pl}/2$. Here, the asymptotically safe value lies squarely in the Type I region, reinforcing the conclusion that asymptotic safety disfavors Type II geometries.

These diagrams suggest that within the constraints of a UV-complete trajectory, only a limited class of black hole mimickers survives, depending on the gravitational parameters $\{M,\ S_2^-\}$ of the object.

7 Discussion

This chapter rounds out the analysis by reflecting on what has been achieved, what its implications are, and what limitations or open questions remain. While much of the work done in this thesis can be seen as a proof of concept rather than a definitive physical prediction, it still offers a concrete case study of how FRG methods and, hence, the asymptotic safety framework can be applied to make predictions in gravitational settings.

7.1 Physical Implications

Let us start with what the results suggest, if taken at face value. We found that a unique UV-complete RG trajectory connects the UV-completing NGFP to the GFP, which fixes the Wilson coefficient of the Weyl-squared term G_{C^2} in the IR. This constrains the phase diagram of gravitational solutions to those compatible with this unique RG trajectory — and excludes other regions of the phase diagram.

However, the interpretation of the physical meaning of these results is not immediately clear. The ADM mass M is no longer just the mass of the black hole, as terms in the expansion explored in Section 6.5 have been reabsorbed into the mass. Furthermore, there is no physical intuition (yet) behind the Yukawa charge S_2^-, even though it is a parameter which distinguishes a black hole mimicker. What then does S_{-2} represent physically? It seems to encode information about the gravitational field's short-distance structure, but whether it corresponds to a real, observable charge is speculative. These issues do not invalidate the RG results, but they do complicate their interpretation.

There is also the question of the *magnitude* of the corrections. The quantum gravitational effects extracted here become significant only at or near the Planck scale. Observing Planckian black holes, or even Planck-scale corrections to astrophysical black holes, is far beyond the reach of current (or foreseeable) technology in observational astrophysics. While this limits the short-term observability of these predictions, it does not lessen their conceptual importance.

Importantly, this work is grounded in a truncation — we have kept only a limited subset of possible couplings. Nonetheless, as a thesis concerned primarily with demonstrating a machinery and not necessarily with the final form of a "correct" theory of quantum gravity, it is valuable to ask: if the Einstein-Weyl truncation were the "right one", what would it imply about the structure of gravity — and how we would we specifically observe it as black hole and black hole mimickers?

7.2 Challenges

Having discussed the physical interpretation of our analysis, it is mandatory to discuss some of its drawbacks and challenges.

First of all, most of the FRG computations are done in Euclidean signature even though gravity is a Lorentzian theory. The Wick rotation between these leaves the calculations much more manageable, but at the expense of some physical principles and predictability — although work has been done in computing asymptotic safety amplitudes using a Lorentzian signature [11].

As a second drawback, while certain combinations of Wilson coefficients are gauge-invariant and scheme-independent quantities, approximations made in the FRG analysis may slightly impact their precise values. Nonetheless, it is expected that the qualitative features of our analysis, and in particular the sign and order of magnitude of G_{C^2}, will stay the same in extended truncations.

Overall, our analysis is situated in a larger "landscape" of challenges faced by quantum gravity research. The most obvious critique of all things quantum gravity is the lack of experimental access. For example, the corrections found here to black hole structure — though interesting and conceptually important — occur at the Planck scale, which remains far out of reach observationally.

From the side of ASQG, and in a broader context, a deeper question emerges: is it fundamental, or is it yet another effective description of some deeper UV-complete theory? Whether or not it is fundamental, the asymptotic safety framework provides a powerful conceptual toolbox for discussing the implications of quantum gravity, and could lead to the next big step in understanding quantum gravity — or even a grand unified theory.

7.3 Outlook

The natural next step is to go beyond the Einstein-Weyl truncation. Including the full quadratic theory terms as in [8] would allow for a broader exploration of the gravitational theory space. Even more ambitious would be to include cubic curvature terms, which have been shown to resolve the pesky singularities inside Schwarzschild black holes, although with the price of a second horizon[46]. Such extensions would test whether the predictions seen here persist under more complete truncations.

The methods and perspective discussed in the thesis — grounded in the FRG and Wilsonian reasoning — remain valid and useful regardless of the truncation. As such, they offer a flexible and powerful approach to the ongoing challenge of understanding quantum gravity.

8 Conclusion

Mapping out the set of black hole-like solutions stemming from ultraviolet (UV)-complete theories of quantum gravity remains an outstanding open challenge, but this thesis offers a small yet concrete step forward within the asymptotic safety framework.

We began by outlining some of the motivations for quantizing gravity: the presence of singularities in classical general relativity (GR) and the perturbative non-renormalizability of GR as a quantum field theory (QFT). In response to these challenges, the asymptotic safety scenario offers a compelling and minimalistic proposal: that gravity might be well-defined at all energy scales and non-perturbatively renormalizable if its renormalization group (RG) flow approaches a non-Gaussian — i.e. interacting — fixed point (NGFP) in the UV.

Our work began by summarising the conceptual foundations of the RG and effective field theory (EFT). Through pedagogical analogies like the Ising model and coarse-graining arguments, we developed the intuition and machinery behind how QFTs evolve with scale and how asymptotic safety proposes to tame the UV divergences that plague GR in perturbative QFT. These ideas were formalised using the functional renormalization group (FRG), which allows for the non-perturbative exploration of gravity or other systems in the regime of strong coupling.

We then applied this framework to a specific model: Einstein-Weyl gravity. This is a simple extension of GR that includes a Weyl-squared curvature term. The philosophy is to regard this as a truncation of the full effective action stemming from asymptotic safety. Using beta functions derived from previous work, we identified an NGFP at $(g^*, g^*_{\mathrm{C}^2}) = (1.0053, 0.7277)$, and extracted the unique UV-complete RG trajectory that flows from this fixed point to the Gaussian fixed point (GFP) in the infrared (IR). Along this trajectory, the Wilson coefficient of the Weyl-squared term was then determined to be $G_{\mathrm{C}^2} = 0.5092\, m_{\mathrm{Pl}}^{-2}$. Using this quantum gravity-determined value, we placed constraints on the classical phase diagram of black hole mimickers in Einstein-Weyl gravity. In addition to Schwarzschild black holes, this phase diagram includes exotic objects encoded by the addition of the higher-order Weyl-squared term, like non-Schwarzschild black holes, asymmetric wormholes and naked singularities. But only a subset of this landscape is compatible with asymptotic safety. The determined value

of the Wilson coefficient selects a unique region in phase space — disfavouring, in particular, Bachian singularities. These results provide a concrete example of how quantum gravity can yield constraints on black hole physics and offer a proof of concept for the strength and predictive power of the asymptotic safety framework.

Overall, this thesis highlighted how computing quantum gravity landscapes — within asymptotic safety or beyond — can help us not only test theories at a theoretical level, but also bridge microscopic quantum dynamics and macroscopic spacetime geometry, shedding new light on the nature of quantum black holes.

9 References

[1] Steven Weinberg. “ULTRAVIOLET DIVERGENCES IN QUANTUM THEORIES OF GRAVITATION”. In: *General Relativity: An Einstein Centenary Survey*. 1980, pp. 790–831.

[2] Martin Reuter and Frank Saueressig. *Quantum Gravity and the Functional Renormalization Group: The Road towards Asymptotic Safety*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 2019.

[3] N. Dupuis, L. Canet, A. Eichhorn, W. Metzner, J. M. Pawlowski, M. Tissier, and N. Wschebor. “The nonperturbative functional renormalization group and its applications”. In: *Phys. Rept.* 910 (2021), pp. 1–114. arXiv: 2006.04853 [cond-mat.stat-mech].

[4] Ivano Basile, Luca Buoninfante, Francesco Di Filippo, Benjamin Knorr, Alessia Platania, and Anna Tokareva. *Lectures in Quantum Gravity*. 2025. arXiv: 2412.08690 [hep-th].

[5] M. Reuter. “Nonperturbative evolution equation for quantum gravity”. In: *Phys. Rev. D* 57 (1998), pp. 971–985. arXiv: hep-th/9605030.

[6] Wataru Souma. “Non-Trivial Ultraviolet Fixed Point in Quantum Gravity”. In: *Progress of Theoretical Physics* 102.1 (July 1999), pp. 181–195.

[7] M. Reuter and Frank Saueressig. “Renormalization group flow of quantum gravity in the Einstein-Hilbert truncation”. In: *Phys. Rev. D* 65 (2002), p. 065016. arXiv: hep-th/0110054.

[8] Benjamin Knorr. *The derivative expansion in asymptotically safe quantum gravity: general setup and quartic order*. 2021. arXiv: 2104.11336 [hep-th].

[9] Astrid Eichhorn and Marc Schiffer. “Asymptotic safety of gravity with matter”. In: (Dec. 2022). arXiv: 2212.07456 [hep-th].

[10] Mikhail Shaposhnikov and Christof Wetterich. “Asymptotic safety of gravity and the Higgs boson mass”. In: *Phys. Lett. B* 683 (2010), pp. 196–200. arXiv: 0912.0208 [hep-th].

[11] Álvaro Pastor-Gutiérrez, Jan M. Pawlowski, Manuel Reichert, and Giacomo Ruisi. $e^+e^- \to \mu^+\mu^-$ *in the Asymptotically Safe Standard Model*. 2024. arXiv: 2412.13800 [hep-ph].

[12] Alessia Platania. “Black Holes in Asymptotically Safe Gravity”. In: *Handbook of Quantum Gravity*. Springer Nature Singapore, 2023, pp. 1–65.

[13] Raúl Carballo-Rubio, Francesco Di Filippo, Stefano Liberati, *et al. Towards a Non-singular Paradigm of Black Hole Physics*. 2025. arXiv: 2501.05505 [gr-qc].

[14] David Tong. *Statistical Field Theory*. http://www.damtp.cam.ac.uk/user/tong/sft.html. Lecture notes, University of Cambridge, Part III Mathematical Tripos. 2017.

[15] S. Silveravalle and A. Zuccotti. “Phase diagram of Einstein-Weyl gravity”. In: *Physical Review D* 107.6 (Mar. 2023).

[16] Matthew D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, 2013.

[17] Luca Buoninfante *et al.* “Visions in Quantum Gravity”. In: (Dec. 2024). arXiv: 2412.08696 [hep-th].

[18] Joseph Polchinski. “Renormalization and Effective Lagrangians”. In: *Nucl. Phys. B* 231 (1984), pp. 269–295.

[19] Christof Wetterich. “Exact evolution equation for the effective potential”. In: *Phys. Lett. B* 301 (1993), pp. 90–94. arXiv: 1710.05815 [hep-th].

[20] Ivano Basile and Alessia Platania. “Asymptotic Safety: Swampland or Wonderland?” In: *Universe* 7.10 (Oct. 2021), p. 389.

[21] Kenneth G. Wilson. “Renormalization group and critical phenomena. 1. Renormalization group and the Kadanoff scaling picture”. In: *Phys. Rev. B* 4 (1971), pp. 3174–3183.

[22] Kenneth G. Wilson. “Renormalization group and critical phenomena. 2. Phase space cell analysis of critical behavior”. In: *Phys. Rev. B* 4 (1971), pp. 3184–3205.

[23] K. G. Wilson and John B. Kogut. “The Renormalization group and the epsilon expansion”. In: *Phys. Rept.* 12 (1974), pp. 75–199.

[24] Kevin Falls, Daniel F. Litim, Konstantinos Nikolakopoulos, and Christoph Rahmede. “Further evidence for asymptotic safety of quantum gravity”. In: *Phys. Rev. D* 93.10 (2016), p. 104022. arXiv: 1410.4815 [hep-th].

[25] Kevin Falls, Callum R. King, Daniel F. Litim, Kostas Nikolakopoulos, and Christoph Rahmede. “Asymptotic safety of quantum gravity beyond Ricci scalars”. In: *Phys. Rev. D* 97.8 (2018), p. 086006. arXiv: 1801.00162 [hep-th].

[26] Alessio Baldazzi, Kevin Falls, Yannick Kluth, and Benjamin Knorr. “Robustness of the derivative expansion in Asymptotic Safety”. In: (Dec. 2023). arXiv: 2312.03831 [hep-th].

[27] Astrid Eichhorn and Aaron Held. “Top mass from asymptotic safety”. In: *Phys. Lett. B* 777 (2018), pp. 217–221. arXiv: 1707.01107 [hep-th].

[28] Jannik Fehre, Daniel F. Litim, Jan M. Pawlowski, and Manuel Reichert. “Lorentzian Quantum Gravity and the Graviton Spectral Function”. In: *Phys. Rev. Lett.* 130.8 (2023), p. 081501. arXiv: 2111.13232 [hep-th].

[29] Alessia Platania and Christof Wetterich. “Non-perturbative unitarity and fictitious ghosts in quantum gravity”. In: *Phys. Lett. B* 811 (2020), p. 135911. arXiv: 2009.06637 [hep-th].

[30] Benjamin Knorr and Marc Schiffer. “Non-Perturbative Propagators in Quantum Gravity”. In: *Universe* 7.7 (2021), p. 216. arXiv: 2105.04566 [hep-th].

[31] Eran Palti. “The swampland and string theory”. In: *Contemp. Phys.* 62.3 (2022), pp. 165–179.

[32] Benjamin Knorr and Alessia Platania. “Unearthing the intersections: positivity bounds, weak gravity conjecture, and asymptotic safety landscapes from photon-graviton flows”. In: *JHEP* 03 (2025), p. 003. arXiv: 2405.08860 [hep-th].

[33] Gustavo P. de Brito, Benjamin Knorr, and Marc Schiffer. “On the weak-gravity bound for a shift-symmetric scalar field”. In: *Phys. Rev. D* 108.2 (2023), p. 026004. arXiv: 2302.10989 [hep-th].

[34] Junsei Tokuda. “Extension of positivity bounds to non-local theories: IR obstructions to Lorentz invariant UV completions”. In: *JHEP* 05 (2019), p. 216. arXiv: 1902.10039 [hep-th].

[35] Astrid Eichhorn, Marc Schiffer, and Andreas Odgaard Pedersen. “Application of positivity bounds in asymptotically safe gravity”. In: *Eur. Phys. J. C* 85.7 (2025), p. 733. arXiv: 2405.08862 [hep-th].

[36] Clifford Cheung, Junyu Liu, and Grant N. Remmen. “Entropy Bounds on Effective Field Theory from Rotating Dyonic Black Holes”. In: *Phys. Rev. D* 100.4 (2019), p. 046003. arXiv: 1903.09156 [hep-th].

[37] K. S. Stelle. “Renormalization of Higher Derivative Quantum Gravity”. In: *Phys. Rev. D* 16 (1977), pp. 953–969.

[38] Alun Perkins. “Static spherically symmetric solutions in higher derivative gravity”. PhD thesis. Imperial Coll., London, Sept. 2016.

[39] Alfio Bonanno and Samuele Silveravalle. “Characterizing black hole metrics in quadratic gravity”. In: *Phys. Rev. D* 99.10 (2019), p. 101501. arXiv: 1903.08759 [gr-qc].

[40] H. Lü, A. Perkins, C. N. Pope, and K. S. Stelle. “Black Holes in Higher Derivative Gravity”. In: *Phys. Rev. Lett.* 114 (17 Apr. 2015), p. 171601.

[41] H. Lü, A. Perkins, C. N. Pope, and K. S. Stelle. “Spherically Symmetric Solutions in Higher-Derivative Gravity”. In: *Phys. Rev. D* 92.12 (2015), p. 124019. arXiv: 1508.00010 [hep-th].

[42] H. Lü, A. Perkins, C. N. Pope, and K. S. Stelle. “Black holes in D = 4 higher-derivative gravity”. In: *Int. J. Mod. Phys. A* 30.2829 (2015), p. 1545016.

[43] Giorgos Anastasiou and Rodrigo Olea. “From conformal to Einstein gravity”. In: *Physical Review D* 94.8 (Oct. 2016).

[44] M. S. Morris and K. S. Thorne. “Wormholes in space-time and their use for interstellar travel: A tool for teaching general relativity”. In: *Am. J. Phys.* 56 (1988), pp. 395–412.

[45] *How to Control the Precision and Accuracy of Numerical Results*. https://reference.wolfram.com/language/howto/ControlThePrecisionAndAccuracyOfNumericalResults.html. Accessed: 2025-05-19.

[46] Bob Holdom. “On the fate of singularities and horizons in higher derivative gravity”. In: *Phys. Rev. D* 66 (2002), p. 084010. arXiv: hep-th/0206219.