

Electronic Thesis and Dissertation Repository

4-22-2022 10:00 AM

Effective Field Theories: A Philosophical Appraisal

Dimitrios Athanasiou, *The University of Western Ontario*

Supervisor: Smeenk, Christopher, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Philosophy

© Dimitrios Athanasiou 2022

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>



Part of the [Philosophy of Science Commons](#)

Recommended Citation

Athanasiou, Dimitrios, "Effective Field Theories: A Philosophical Appraisal" (2022). *Electronic Thesis and Dissertation Repository*. 8493.

<https://ir.lib.uwo.ca/etd/8493>

This Dissertation/Thesis is brought to you for free and open access by Scholarship@Western. It has been accepted for inclusion in Electronic Thesis and Dissertation Repository by an authorized administrator of Scholarship@Western. For more information, please contact wlsadmin@uwo.ca.

Abstract

The word “effective” has become the standard label attached to scientific theories these days. An effective theory allows us to make accurate predictions about a physical system at a certain (energy, length) scale while being largely ignorant of the details at more fundamental levels. One does not need to know anything about the deeper, quantum structure of water molecules to describe the macroscopic behaviour of waves or water in a glass. Although effective descriptions so broadly construed have been part of research in physics since the earliest stages of modern science, it is particle physics that has most clearly relied on and brought to the fore some of the most interesting and admittedly puzzling aspects of this way of looking at theories. Indeed, the effective field theory (EFT) program in QFT has established itself as the most natural way to understand renormalisation and dissipate initial reservations about the status of these techniques by treating higher-order processes as contributions suppressed at lower energy scales. QFT is thus treated as the “effective” framework par excellence with the decoupling of scales constituting its permeating tenet. The goal of this project is to attempt a philosophical appraisal of EFTs as currently used in high energy physics as well as assess the possibility that the whole program eventually breaks down, i.e. fails to apply when certain preconditions do not hold. Accordingly, the dissertation is logically divided into two parts with the first two chapters dedicated to discussion of the relation between EFTs and traditional questions in the philosophy of science concerning the structure of scientific theories, the formulation and defence of scientific realism as well as its connection to possible ontological readings of EFTs. The second part constitutes an analysis of two well-known problems that have been accorded the status of crises in the physics literature: the hierarchy problem and the cosmological constant problem. Our main focus will be to uncover those assumptions responsible for undermining the validity of the EFT techniques in their respective context. In light of this analysis, we will ultimately lean towards a more cautionary or “reserved” approach to EFTs.

Key words: philosophy, physics, scientific theories, realism, effective realism, quantum field theory, renormalization group, effective field theories, standard model, beyond the standard model, naturalness, fine-tuning, hierarchy problem, cosmological constant, reductionism, theory construction, scientific methodology

Lay Summary

An effective theory allows us to make predictions about a physical system without possessing a complete description of all its microconstituents or without fully tracking their intricate interactions. For example, when we apply Newtonian mechanics in high school physics problems involving balls colliding with one another we typically abstract from their atomic and molecular structure and disregard negligible influences such as interactions with the surrounding air molecules. Including only a handful of input parameters for our models allows us to keep our equations tractable while extracting information up to a desired accuracy. Although effective theories understood as such approximating tools have been part of research in physics since its earliest stages of development (one just recall early applications of Newton's theory of gravity to planetary motions], it is particle physics that has brought to the fore the most interesting and admittedly puzzling aspects of this approach to seeing theories. Perhaps the greatest conceptual insight obtained by an effective reading of QFT is that it was possible to put on a firmer physical ground the process of renormalisation, which, although indispensable in producing meaningful QFT calculations, was typically seen as a trick for sweeping the pathogenies of the theory under the rug. The advent of EFTs changed that. Unsurprisingly, the centrality of this new conception of understanding physical theories calls for key revisions to the way philosophers have traditionally thought about issues such as how is a scientific theory structured, how it represents entities in the world and how it is to be interpreted. One goal of this thesis is to contribute to this "revisionary" project. Another goal is to examine possible limitations of the effective framework, i.e. cases where its tools lead us to incorrect expectations. To this effect, two infamous open problems of modern physics, the hierarchy and cosmological constant problems, are examined in the latter half of the thesis. It is argued that they are signals of a breakdown of EFTs and are used as the source of extracting preconditions for the applicability of the effective framework.

Contents

Abstract	i
Lay Summary	ii
Abbreviations	vi
List of Figures	vi
List of Tables	vii
Acknowledgments	ix
0 Intrata: Physics Background	1
0.1 Quantum Field Theory	1
0.2 Renormalisation	6
0.2.1 Regularisation Schemes	7
0.2.2 Renormalisation Schemes	8
0.3 Renormalisable Theories	9
0.3.1 ϕ^4	10
0.3.2 QED	12
0.3.3 QCD	17
0.4 Renormalisation Group	19
0.4.1 Continuum Version	19
0.4.2 Wilsonian Version	22
0.4.3 Connection to Renormalisation	25
0.5 Effective Field Theory	25
0.5.1 Top-Down EFTs	26
0.5.2 Bottom-Up	28
0.5.3 New Fields	30
1 EFTs as “Theories”	31
1.1 Standard Accounts of Scientific Theories	33
1.1.1 Formal Approaches	33

1.1.2	Informal Approaches	39
1.2	A Plastic View	43
1.2.1	First Level: Stratification	43
1.2.2	Second Level: Strategising	45
1.2.3	Third Level: Combining	48
1.3	The Role of EFTs	50
1.4	What is QFT?	55
1.4.1	The Fraser - Wallace Debate	57
1.4.2	The Status of the Cut-Off	64
1.5	Quick Application: Scientific Representation	67
2	EFTs as Theories of the World	69
2.1	Realism Matters	70
2.2	Effective Realism	75
2.2.1	Motivation: Realism and Pessimism	75
2.2.2	Effective Realism: Definition	77
2.2.3	Effective Realism: Prospects and Limitations	80
2.3	Against (Some) Fundamentalism	85
2.3.1	First Route: Defining Commitments	86
2.3.2	Second Route: Reject Fundamentalism	89
2.3.3	Classical Analogues	96
2.4	How Realism Matters	101
2.4.1	EFTs: An Application... or Three	108
2.5	Conclusions	112
2b	Interlude: The World according to EFTs	114
2b.1	The (Preliminary) Verdict	121
3	When EFTs Seem to Fail I: The Hierarchy Problem	123
3.1	A Naturalness Story	124
3.2	The Hierarchy Problem	130
3.2.1	Formulations	133
3.2.2	Reactions	140
3.3	Reaction 1: New Physics*	142
3.3.1	Prediction or Reconstruction?	150
3.4	Reaction 2: Denial	157
3.4.1	Technical Objections	157
3.4.2	Disputing Fine-Tuning	162
3.4.3	Seeking Motivation	168
3.5	Reaction 3: From Effective to Ineffective Theories?	172
3.5.1	A Failure of Reductionism?	175

3.5.2	Quantum and Classical Analogies	180
3.5.3	Naturalness \subset Autonomy	186
3.5.4	Restrict Domain of Applicability	191
3.5.5	Beyond a Reductionistic Physics?	194
3.6	Conclusions	197
4	When EFTs Seem to Fail II: Cosmology	201
4.1	The Void, the Vacuum and Nothingness	202
4.1.1	The Abhorred Vacuum	202
4.1.2	The Infamous Constant	205
4.2	What is the Problem with the Cosmological Constant?	206
4.2.1	Elastic Spacetime	208
4.2.2	When the Nothing Nothings	210
4.2.3	Turning on the Vacuum... Cleaner	213
4.2.4	The (Real) Problem	218
4.3	EFTs under Pressure...again!	219
4.3.1	Reactions	219
4.3.2	Preconditions for EFTs	220
4.4	Conclusions	226
4b	Interlude: The World beyond EFTs	228
4b.1	Appealing the Verdict	229
4b.2	A (Final) Verdict	230
5	Coda: Philosophical Lessons	232
5.1	Ontology: Realism and Reduction	232
5.1.1	Theory Interpretation in Light of EFTs	236
5.2	Methodology I: Theory Construction	240
5.2.1	Theory Progression with the Help of EFTs	243
5.3	Methodology II: Extra-empirical Guiding Principles	245
	Bibliography	251

Abbreviations

AQFT = Algebraic Quantum Field Theory

BSM = Beyond the Standard Model

EFT = Effective Field Theory

EM = Electromagnetism

ER = Effective Realism

LQFT = Lagrangian Quantum Field Theory

LQG = Loop Quantum Gravity

NRQM = Non-Relativistic Quantum Mechanics

SM = Standard Model.

SR = Special Relativity

SR = Scientific Realism*

ST = String Theory

QCD = Quantum Chromodynamics

QED = Quantum Electrodynamics

QM = Quantum Mechanics

QFT = Quantum Field Theory

QG = Quantum Gravity

List of Figures

1	The basic QFT “ingredients”	5
2	Triangle Diagram	6
3	Four-point scattering for ϕ^4	11
4	Counterterms for ϕ^4	12
5	Photon Propagator Correction	13
6	Electron Propagator Correction	14
7	Vertex Correction	15
8	Higher-order Diagrams	16
9	Loop Diagrams in QCD	18
10	Renormalization Group Trajectories	22
11	Heavy field ψ correction to light field ϕ propagation	27
12	From Electroweak Interaction to Fermi Theory	28
1.1	Euler-Heisenberg Effective Theory	53
2.1	Propagator to Point Interaction	108
3.1	Schematic Representation of the Hierarchy Problem	131
3.2	Fermionic Loop in Scalar Field Propagator	132
3.3	Compactification: the additional dimension of the superspace is “winded” around the subspace of the effective theory	143
3.4	SYSY Cancelling Loops	147
3.5	The evolving conception of naturalness	150
3.6	Non-commutative Field Theory Loop Corrections	194
4.1	The Cosmological Constant Problem: A Decision Tree	208
5.1	The Dialectic of Theory Construction	243

List of Tables

1.1	QFT: AQFT or EFT?	59
1.2	QFT: Both AQFT and EFT	62
2.1	A Realism Typology	104
2b.1	Attitudes to EFTs	114
2b.2	Ontological Temperaments	121
3.1	Naturalness Record	155
3.2	Wilsonian QFT & Classical Statistical Mechanics	182
3.3	The Decoupler's Game	198
3.4	Solving the Decoupler's Game	199

Acknowledgments

Contrary to appearances, a dissertation is not a solitary work. I owe much to many people I have interacted with over the years for the ideas developed in this thesis. First and foremost, I owe a huge amount of gratitude to my supervisor, Christopher Smeenk, for his constant mental, moral and material support. Chris was the instigator of my interest in effective field theories and naturalness problems. He has been most encouraging and patient with all the half-baked ideas I served him throughout the past three years, always probing and supporting me in the pursuit of my intellectual predilections. Our collaboration could not have been smoother during the “trying” times of the global pandemic and the isolation and disruptiveness it entailed. After each of our online meetings, I felt energised and reassured that my project was worth pursuing and interesting ideas could be developed. Needless to add, the way I have come about the issues discussed in this dissertation has been deeply influenced by his comments and overall approach to philosophy of physics.

I am also deeply thankful to have worked with two other major figures in the field of philosophy of science during my years at Western. I am very grateful to Wayne Myrvold for all the philosophy of physics he has taught me, especially on quantum mechanics, statistical mechanics and probability theory. His ideas have greatly influenced my approach to issues of probability in this thesis. His comments on my prospectus defence helped me make my ideas more focused and precise. Wayne’s strive towards clarity and soundness have inspired since my first days at Western – although I may have only asymptotically been able to emulate it.

Last, but not least, I am indebted to Robert DiSalle for all the time he so generously devoted to discussing science and philosophy with me. Our conversations on the structure of scientific theories, logical positivism, Stein and realism have had a profound effect on me. It is not exaggeration to say that half of this project, chapters one and two, would not have been completed, had it not been for the wintery meetings in his cosy, dimly lighted office at Stevenson hall. His often understatedly and humbly presented ideas greatly influenced the way I think about a lot of these issues. Of course, what is presented here is the outcome of the way I understood and “appropriated them” for my purposes. I, therefore, remain solely accountable for any inaccuracies and mistakes.

For the past two years, I was lucky enough to be part of Chris' research group that consistently met online during the pandemic. Keeping in touch with people working in the same field helped me maintain a steady pace in my work while also allowing me to receive feedback and motivation from friends and colleagues – even only through a computer screen. I am thankful to Mark Baker, Craig Fox, Marie Gueguen, Niels Linnemann, Yichen Luo, Matthew Parker for creating a friendly atmosphere, sharing their work with me and offering comments on mine. Special thanks to Marie Gueguen and Niels Linnemann for the extensive comments and the extra time they took to discuss parts of my thesis in separate meetings. A warm thank you to Craig Fox for giving me some encouraging words when it mattered. Finally, a big thank you to Adam Koberinski for his extensive comments on earlier drafts of my work, for extra discussions on the philosophy of QFT and his work on the cosmological constant as well as his reading suggestions.

Before the pandemic, I was also very lucky to partake into the wider philosophy community at Western. There are several people that, in their own way, contributed to the completion of this thesis. I am most grateful to Todd Nagel and Martin Zelko for spending hundreds of hours with me, both in person and online, delving into general relativity, the foundations of quantum mechanics, statistical mechanics and the history of spacetime theories as part of our weekly reading group meetings. I am particularly grateful to friend, colleague and roommate Martin for all the discussions we have had on philosophy, physics, Quine and our everyday troubles. I have learned so much from him during the past four years. Special thanks are also due to two people that made even the snowiest of days in London feel as warm as home: John Lehmann and Ashkan Parchizadeh. I am grateful to Ashkan for the countless hours we spent together on and off campus having fun, discussing history, philosophy and politics as well as exchanging cultural perspectives on Greece and Iran. Truly, life would have been so much harder without Ashkan's help. I am grateful to John for all our deep conversations on math, philosophy and life over the smell coffee – and, occasionally, beer. His friendship was a source of support during crucial moments of self-doubt and disappointment. As I have told him, sharing my existential burdens with him has always had a cathartic effect on me. Many other people deserve my gratitude for the time I spent in Canada but special thanks are due to: Tom de Saegher, Sona Gosh, Phillipos Papagiannopoulos and Greg Lee.

Owing to the pandemic, the project became a transatlantic affair as most of it was meant to be completed back home, in Greece. The last couple of years have been truly trying psychologically and it would have been absolutely impossible to pull this huge endeavour through, had I not been blessed with some of the deepest and longest friendships. I am forever grateful to Tonia Thoma for being there whenever it mattered, Constantina Fotopoulou for being my psychotherapist during the hardest of times and Giorgos Goulas for being, well, Giorgos – ever

since we took the GRE exams. My most sincere token of gratitude is due to Haris Tzerefos for having put up with my tantrums, despair and craziness for the past (too many) months. Thank you for me teaching QFT for the first time a few years ago and for all the time we spent studying and discussing physics, philosophy, history, psychology, poetry, my thesis and life. Thank you, guys! It is impossible to even imagine what life would be like without you. Many thanks to the remaining members of the “Order of the Faithful Dog”: Afroditi, Dimitris and Angeliki, Dimitris, Marilina and Ioanna for all that we have been through over the years.

Finally, I owe the utmost gratitude to my mother for all that she has done for me over the course of a lifetime and especially during some of the hardest years of her life – coinciding with my PhD. There might not be many constants in this world, but her unwavering love and support for me had better count as one. She had to suffer through a lot for me to be writing these lines now. In many respects, this thesis is more her work than mine. It is only appropriate that it be dedicated to her.

Final note: This project has been generously funded by the Templeton Foundation New Directions in Philosophy of Cosmology grant and the Rotman Institute of Philosophy at Western. I am very much indebted to the Canadian state and people for offering me the chance to study at one of their best institutes in the field of philosophy through their taxes and good will.

Chapter 0

Intrata: Physics Background

Here we introduce some of the basic physics tools that will be needed for the rest of the discussion. Most of the main physical results alluded to later in the text are collected here for ease of access and to preserve the logical order of the presentation¹.

0.1 Quantum Field Theory

The quantum theory of fields (or QFT for short) is one of the main pillars of contemporary physics. It is the framework currently employed to describe three of the four known fundamental interactions in nature. Electromagnetic and nuclear phenomena at the smallest distance scale ever probed by man are described in stunning accuracy using the arsenal of QFT. Effective field theories themselves might be said to only comprise a certain way of thinking about QFT. It is thus important to familiarise ourselves with the basic concepts of this theoretical framework.

Quantum Fields QFT is a theory that aspires to combine quantum mechanics (QM) with special relativity (SR). After early attempts to tackle this problem directly failed (e.g. Klein-Gordon equation), it became clear that the only consistent way to preserve the causal structure of relativity with the weirdness of quantum fluctuations and the uncertainty relations is to turn to a field-theoretic approach.

How does one go about writing down quantum fields, however? Typically, the idea is to start with some classical field theory like the Lagrangian of the electromagnetic field and attempt to quantise it, i.e. subject it to a process whereby fields are upgraded to operators and appropriate

¹Although I consulted multiple textbooks to create this short summary (references are all over the text), my notation, for the most part follows that of Schwartz 2014.

commutation relations are imposed. This is to mimic the process of quantisation of position and momentum in the transition from classical to quantum mechanics where the Poisson bracket is mapped to the appropriate commutation relation: $\{A, B\} \rightarrow [A, B]$. This procedure, which is called canonical quantisation, leads to conditions of the form²:

$$[\hat{\phi}(t, \mathbf{x}), \hat{\Pi}^0(t, \mathbf{y})] = i\delta^{(3)}(\mathbf{x} - \mathbf{y}) \quad (0.1)$$

where $\phi(x)$ is some operator-valued field and $\Pi(x)$ its conjugate momentum with $x = (t, \mathbf{x})$ the spatiotemporal coordinate vector. One can define a Hamiltonian to describe the evolution of the field. Fields are expanded in terms of creation and annihilation operators, which acting on the vacuum state (the state of lowest) energy describe excitations of the field. As a consequence, a field will be written as an expansion of the form over all possible momenta k :

$$\phi(x) = \int \frac{d^4k}{N(k)} \left[a^\dagger(\mathbf{k})e^{-ikx} + b(\mathbf{k})e^{ikx} \right] \quad (0.2)$$

with $N(k)$ some renormalisation constant that need not worry us here. Note that, in general, the a and b do not have to be the conjugate of one another - this only happens when the quantum (excitation) of the field is its own antiparticle. More complex fields such as spinor or vector fields will also need to be expanded in appropriate bases to take into account the polarisation of the described particle. This means that field is expanded over a further basis representing the possible polarisations of the particle:

$$\psi_\alpha(x) = \int \frac{d^4k}{N(k)} \left[\sum_s a^{s\dagger} u_\alpha^s(k)(\mathbf{k})e^{-ikx} + b^s(\mathbf{k})u_\alpha^s(k)e^{ikx} \right] \quad (0.3)$$

Here s could be labelling the possible spin values $+1/2, -1/2$ and α the components of the spinor field $\psi(x) = (\psi_1(x), \psi_2(x), \psi_3(x), \psi_4(x))$.

Path Integral Another powerful tool that can be used in QFT is the so-called path integral, which was first devised by Dirac and then exploited by Feynman to produce an equivalent formulation for QM. In fact, the path integral is the most frequently tool used in QFT because it sidesteps the complexities of the canonical quantisation picture – which become particularly

²A note on notation: strictly speaking, since fields are operator-valued fields, they must carry a hat at all times to indicate this – as is the case in the following equation. For convenience, we are dropping hats in the rest of the text.

pressing when dealing with gauge theories.

The idea behind the path integral formulation is to sum over all possible field configurations that can possibly exist between two spacetime points. Much like in Lagrangian approaches to mechanics, one starts with the action of the field theory $S[\phi]$ which returns the n -point functions of the theory once we sum (integrate) over all possible configurations:

$$G(x_1, x_2, \dots, x_n) = \langle 0 | T[\hat{\phi}(x_1) \dots \hat{\phi}(x_n)] | 0 \rangle = \int \mathcal{D}\phi e^{iS[\phi]} \phi(x_1) \dots \phi(x_n) \quad (0.4)$$

To make this happen, we construct the so-called generating functional $\mathcal{Z}[J]$ which encapsulates all the relevant information for our theory:

$$Z[J] = \int D\phi e^{i \int d^4x [\mathcal{L}_0[\phi] + \mathcal{L}_{int}[\phi] + J(x)\phi(x)]} \quad (0.5)$$

where $J(x)$ are called source terms. This quantity can then be used to systematically extract propagators and n -point functions by taking derivatives with respect to the source term at specific points in spacetime. For example, one calculate the propagator for a field to evolve from point x to point y :

$$\langle 0 | T[\hat{\phi}(x)\hat{\phi}(y)] | 0 \rangle = (-i)^2 \frac{1}{Z[0]} \frac{\delta^2 Z[J]}{\partial J(x) \delta J(y)} \Big|_{J=0} \quad (0.6)$$

The source terms can be thought of as ‘‘sourcing’’ or giving rise to the field excitations at spacetime points. If one takes more complicated derivatives (e.g. at more than two spacetime points) they can unpack more and more information out of the generating functional. With enough mathematical care (using Grassmann variables), the functional integral can easily be used for fermionic fields as well (e.g. see Srednicki 2007, ch. 44).

Evidently, the path integral approach is a very powerful tool that allows us to extract systematic information for any field action we write down. This comes particularly handy in the treatment of gauge theories where gauge conditions can be imposed as restrictions on the integral. This is not all, however. A very important feature is that we can modify the ‘‘action input’’ accordingly to narrow the scope of the information we want to extract:

- $\mathcal{Z}[J]$: the full generating functional returns all the possible n -point functions for a given action S input
- $\mathcal{W}[J]$: the functional for connected diagrams [$Z[J] = e^{iW(J)}$] only returns those n -point

functions that do not consist of two separate, unconnected diagrams; so-called bubble diagrams are also excluded

- $\Gamma[A]^3$: performing a Legendre transformation on \mathcal{W} constitutes a further restriction and only returns the so-called 1PI diagrams, i.e. those that cannot be split into two distinct diagrams when an internal line is severed

The third functional, also-called the effective action, is particularly useful in discussions of renormalisation because there, one wants to focus on the real “troublemakers”, i.e. the loop diagrams that are responsible for the appearance of infinities.

Interactions This is all good for free fields, but nature requires more of us than just free fields. One must also be able to describe the way fields interact with one another. In QFT a full analytical solution is not available and we have to rely on some kind of approximation techniques. Perturbation theory, in particular, plays an absolutely crucial role in standard developments of the subject. To treat an interaction perturbatively, it is assumed that the Hamiltonian of the interacting system can be split into two parts:

$$H = H_0 + H_{int} \quad (0.7)$$

with H_0 representing the free theory and H_{int} a small perturbation on that. The end goal is to calculate vacuum expectation value (VEVs), that is expressions of the form:

$$\langle O | T \phi(x_1) \phi(x_2) \dots \phi(x_n) | 0 \rangle \quad (0.8)$$

where T signals that the product has been put in appropriate time ordering to avoid any ambiguities in integration. A lot of technical machinery is then constructed in order to keep track of this rather complicated task. The key idea is to treat interactions as “interrupting” the propagation of fields at specific points in space-time. Feynman’s ingenious representation of the processes in terms of diagrams is then employed to keep track of the calculations. The basic ingredients are (diagram adjusted from Peskin and Schroeder 1995 (1995, p. 94):

a) **propagators**: representing the propagation of a field from spacetime point x to y

b) **vertices**: representing points where fields meet, i.e. interact

³For the sake of completeness, this defined as: $\Gamma(A) = - \int dx JA + W(J)$ where $A = \frac{\partial W(J)}{\partial J}$ and $J = \frac{\delta \Gamma(A)}{\delta A}$

c) **(external) points**: representing excitations that can be detected




Figure 1: The basic QFT “ingredients”

Processes such as 0.8 will need to be computed in order to estimate the value of certain interaction amplitudes:

$$\mathcal{A} = \langle q_1, \dots, q_m | \hat{S} | p_1, \dots, p_n \rangle \quad (0.9)$$

where $p_i, i = 1, \dots, n$ represent incoming particles and $q_j, j = 1, \dots, m$ outgoing particles. The matrix \hat{S} encodes the information required to calculate the amplitude of the (scattering) process and can be written in terms of the quantum fields involved as:

$$\hat{S} = T \left[\exp \left\{ -i \int d^4z \hat{H}_{int} \right\} \right] \quad (0.10)$$

Expanding this out, one obtains expressions involving products of fields at spacetime points such as $\phi(x)\phi(x)\psi(y)\psi(y)$. The in- and out- states themselves are then also written in the form of creation and annihilation operators acting on the vacuum so that we end up with expressions of the form:

$$\langle 0 | a(\mathbf{k})\psi^\dagger(\mathbf{y})\psi(\mathbf{y})\psi^\dagger(\mathbf{x})\psi(\mathbf{x})a^\dagger(\mathbf{k}) | 0 \rangle \quad (0.11)$$

Using a key theorem by Wick it is then possible to reduce the whole problem into a calculation of pairwise propagators like $\Delta_F(x - y) = \langle 0 | T[\phi(x)\phi(y)] | 0 \rangle$. So, any diagram can be broken down to the few basic components we saw in Fig. 1. Feynman diagrams are drawn based on this combinatorics game, their respective integral expressions are estimated and then summed up to obtain the final result for the whole process (see Lancaster and Blundell 2014, sect. 20.1 for a detailed discussion).

0.2 Renormalisation

If the complexity of the subject already seems pretty daunting, the worst is yet to come. For as soon as one attempts to actually calculate probability amplitudes beyond the tree level (i.e. at orders where loop corrections become relevant), they are faced with the bleak prospect of calculating infinite integrals! Diagrams such as diagram 2 (taken from wikimedia commons).




Figure 2: Triangle Diagram

lead to infinite amplitudes. To fix this nonsensical result one subjects the theory to a process of renormalisation whereby infinities arising at graphs of high orders are “re-absorbed” into the parameters of theory. This is allowed because we can reparametrise the Lagrangian, i.e. shift the fields, coupling constants and masses appearing in it. Usually we present this shift as a splitting of these terms into two parts:

$$\mathcal{L}_R = \mathcal{L}_B + \mathcal{L}_{ct} \quad (0.12)$$

Once we run the calculations of problematic diagrams again, we can adjust the second part so as to cancel the divergence and result in a finite value. This algorithmic procedure consists of roughly the following steps (also consult Folland 2013 ch. 7):

1. Re-parametrise the Lagrangian so that all the parameters can be adjusted later on. Usually, one splits them into two parts: a “bare”, which contains the infinity, and a “counterterm”, which kills off the infinity in calculations.
2. Introduce a regulator Λ, ϵ, \dots which will help one separate the infinite quantities in the calculation. Regulators can come in many forms but one can very roughly see them as tools to discretise the theory and avoid taking into account too many degrees of freedom.
3. Impose a renormalisation condition which will help fix the value of the parameters to a specific finite value at a given scale. This is frequently chosen empirically.

4. Send the regulator to infinity (or zero) to return to a continuum theory. There must be no dependence on this unphysical parameter.

0.2.1 Regularisation Schemes

The process of regularisation is a formal technique meant to make manifest the divergence to be eliminated. There are many ways to regularise a theory, but since picking one is an arbitrary choice from a physical (though not from a practical) perspective, they must all lead to the same physics. The type of regulator will also determine the type of limit we will take (to 0 or ∞). Some of the most frequently employed schemes are (see also Schwartz 2014, appendix B):

1. **Lattice:** we essentially discretise spacetime by introducing a grid with some minimal length α for point distances. We then evaluate our fields only over these points instead of the whole continuum. The limit we take is $a \rightarrow 0$. Despite its simplicity this method has the disadvantage of breaking the Lorentz invariance of the theory (the place on the lattice matters).
2. **Hard Cut-off:** we change the structure of the propagator by imposing an upper bound to the possible values of the momentum:

$$G(k) = \begin{cases} \frac{1}{k^2+m^2} & |k| < \Lambda \\ 0 & |k| > \Lambda \end{cases}$$

After all calculations are performed we expect independence from the parameter Λ and take the limit $\Lambda \rightarrow \infty$. This method seems the most intuitive but also breaks Lorentz invariance as it can be seen as equivalent to putting the theory on a lattice (in place of very small distances we disallow very high momenta).

3. **Pauli-Villars Regularisation:** fictitious or ghost particles with large masses M are added for each particle of mass m in the Lagrangian. Their purpose is to cancel with the loop amplitudes of the physical particles at large momenta. This method leads to the following modification of the integrals:

$$\int \frac{d^4k}{(2\pi)^4} \left[\frac{1}{(k^2 - m^2 + i\epsilon)^2} - \frac{1}{(k^2 - M^2 + i\epsilon)^2} \right] = \dots = -\frac{i}{16\pi^2} \log\left(\frac{m^2}{M^2}\right) \quad (0.13)$$

As the mass M becomes large (i.e. as we take the limit $M \rightarrow +\infty$, the ghost particles decouple from all other fields and drop out of the calculations. This approach preserves Lorentz invariance but it can quickly become impractical when multiple ghost fields need to be introduced while it also fails to respect gauge invariance (and thus cannot be used in QCD)⁴.

4. **Dimensional Regularisation:** one makes the, admittedly bizarre, move of changing the dimensions of the theory from D to $d = D - \epsilon$, so that the integral now looks like:

$$\int \frac{d^d k}{(2\pi)^d} \frac{1}{(k^2 - m^2 + i\epsilon)^2} \quad (0.14)$$

This shift to non-integer dimensions has the advantage, as we shall see in more detail, of keeping the gamma functions obtained from blowing up (much like calculating $\frac{1}{x}$ at an $x = \epsilon \neq 0$). Although this approach makes the least sense physically, it is the most widely used in QFT as it leads to the easiest calculations in most domains (like QCD).

Needless to add, each scheme comes with its own advantages and disadvantages with some being more intuitive and others respecting more symmetries of the system. Due to the latter property and the better computational tractability that it offers, dimensional regularisation is the most widely employed scheme of all.

0.2.2 Renormalisation Schemes

As with regularisation so with renormalisation: there is a variety of conditions one can impose to derive finite results. The whole trick amounts to fixing the value of specific quantities such as the pole of a propagator at a given momentum in order to fix the value of the relevant parameter.

- **On-shell Conditions:** perhaps the most physically salient choice is to fix the value of the parameter equal to its physical value. What this value is typically depends on some particular quantity. For example, the counter-term for the mass of the electron is set so that the pole of the propagator:

$$G(\not{p}) = \frac{i}{\not{p} - m_R + \Sigma_R(\not{p})} \quad (0.15)$$

⁴The reason for this is that the introduction of mass for the gauge bosons breaks gauge invariance as e.g. in the case massive electromagnetism.

comes out equal to the physical mass m_P , i.e. $\Sigma_R(/p) = m_R - m_P$ where $\Sigma_R(/p)$ contains the contributions from loop diagrams. In effect, the on-shell condition here says that the finite parts of the counterterms are chosen so as to make m_R come out equal to m_P . We will see that the counterterms for QED are fixed using conditions on the photon and electron propagator as well as the interaction vertex.

- **Minimal Subtraction:** the easiest way to get rid of infinities after regularising the theory is to identify the infinity (or better: the part that will give rise to an infinity after taking the final limit) and directly eliminate it. So, starting with a result of the form:

$$\Sigma = \frac{\alpha}{2\pi} \left[\frac{1}{2} \not{p} \left(\frac{2}{\epsilon} + \ln(4\pi e^{-\gamma_E}) \right) - 2m_R \left(\frac{2}{\epsilon} + \ln(4\pi e^{-\gamma_E}) \right) + \text{finite} \right] \quad (0.16)$$

the counterterm is chosen so as to eliminate terms containing $\frac{1}{\epsilon}$ (or Λ in some other regularisation scheme), i.e. $\delta_2 = \frac{\alpha}{4\pi} \left(\frac{2}{\epsilon} + \ln(4\pi e^{-\gamma_E}) \right)$ and $\delta_m = -\frac{3\alpha}{4\pi} \left(\frac{2}{\epsilon} + \ln(4\pi e^{-\gamma_E}) \right)$. Minimal subtraction is a particularly easy scheme to work with and it is indeed the most widely in QFT - especially for gauge theories.

Any important condition is that the physics stays the same whichever regularisation scheme we choose and whatever renormalisation conditions we impose. This means that we are granted considerable freedom in specifying the finite part that will be preserved after renormalisation. Of course, depending on choices made at the level of the coupling constants adjustments might need to be made to other parameters of the theory such as fields to preserve the invariance of the final amplitude. As we will see, this will be the basis for the idea of the renormalization group.

0.3 Renormalisable Theories

It is obvious that the success of renormalisation is a crucial matter for the treatment of any QFT. Not all theories are well-behaved, however. Depending on their “friendliness” to renormalisation, theories can be classified into: a) super-renormalisable, b) renormalisable and c) non-renormalisable. The distinction is based on the number of parameters that need to be adjusted in order to obtain finite results. For renormalisable or super-renormalisable theories only a finite number of parameters $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$ need to be adjusted so that once this set is fixed, the

theory produces finite answers up to arbitrary orders. Non-renormalisable theories, by contrast, fail in this regard: at each order new parameters that need to be tweaked show up.

Fortunately, as it turns out, the main QFTs comprising the standard model of particle physics are renormalisable. These include QED, QCD, the theory of electroweak interactions and the Higgs mechanism. Proof of these facts by people like t' Hooft and Gross was a major point in favour of the model. Here, we briefly examine the simple case of ϕ^4 theory and QED –only mentioning some basic facts about QCD– to illustrate the main concepts we have been developing in the abstract so far.

0.3.1 ϕ^4

The Lagrangian for ϕ^4 theory is all too familiar:

$$L = \frac{1}{2}(\partial_\mu\phi_B)^2 - \frac{m_B^2}{2}\phi_B^2 - \frac{\lambda_B}{4!}\phi_B^4 \quad (0.17)$$

Instead of working with this bare Lagrangian in which ϕ, m, λ are all bare (unphysical, non-renormalised) parameters, we renormalise them by introducing the following transformation:

$$\begin{aligned} \phi_B &\rightarrow \sqrt{Z}\phi_R \\ m_B &\rightarrow Z_m m_R \\ \lambda_B &\rightarrow Z_\lambda \lambda_R \end{aligned}$$

so that the Lagrangian becomes:

$$L_R = \frac{Z}{2}(\partial_\mu\phi_R)^2 - \frac{ZZ_m^2 m_R^2}{2}\phi_R^2 - \frac{Z_\lambda Z^2}{4!}\lambda_R\phi_R^4 \quad (0.18)$$

Notice that for any parameter, like λ we can now write the bare version as the sum:

$$\lambda_B = \lambda_R + \delta\lambda \quad (0.19)$$

of the renormalised parameter λ_R plus a “correction” term $\delta\lambda$. The correction coefficients correspond to the difference between bare and renormalised parameters and we relate them to the rescaling factors Z_r (they are the rescaling factors just expressed in a different way) as:

$$Z_r = 1 + \delta_r \quad (0.20)$$

With this alternative notation, we can perform the split of the Lagrangian into the two parts we mentioned above: $\mathcal{L}_B + \mathcal{L}_{ct}$ to write:

$$\mathcal{L}_R = \frac{1}{2}(\partial_\mu \phi_B)^2 - \frac{m_B^2}{2}\phi_B^2 - \frac{\lambda_B}{4!}\phi_B^4 + \frac{\delta_\phi}{2}(\partial\phi)^2 - \frac{\delta_m m^2}{2}\phi^2 - \frac{\delta_\lambda}{4!}\lambda\phi^4 \quad (0.21)$$

Let's see how all of this translates for scattering amplitude calculations. In the perturbative expansion for the scattering of two particles up to second order we obtain the following Feynman diagrams (see figure 3 adapted from Lancaster and Blundell 2014, p. 286).




Figure 3: Four-point scattering for ϕ^4

The first diagram is of course first-order and can be evaluated as:

$$iM_a = -i\lambda \quad (0.22)$$

with the rest being second-order and described by the integral:

$$\mathcal{M} = \int \frac{d^4q}{(2\pi)^4} \frac{i}{q^2 - m^2 + i\epsilon} \frac{i}{(p - q)^2 - m^2 + i\epsilon} \quad (0.23)$$

which depending on the form of the interaction depicted by the diagram will lead to some quantity that will depend logarithmically on Λ . Indeed, $\mathcal{M} \propto \ln\left(\frac{\Lambda}{p}\right)$. Obviously, this dependence on Λ is bound to create problems when $\Lambda \rightarrow \infty$, so we will need to adjust the relevant counterterms in Λ_R to cancel the potential infinity. Using a scheme like dimensional regularisation one adds the following terms:

$$\delta_m = \frac{\lambda}{16\pi^2} \frac{1}{\epsilon} m^2 \quad (0.24)$$

$$\delta_\lambda = \mu^\epsilon \frac{\lambda^2}{16\pi^2} \frac{3}{\epsilon} \quad (0.25)$$

The additional parameter μ is added in dimensional regularisation to keep the parameters dimensionless. Translated into the language of Feynman diagrams these new terms correspond to extra diagrams of the form (diagram adjusted from Lancaster and Blundell 2014):



Figure 4: Counterterms for ϕ^4

0.3.2 QED

ϕ^4 is a simple enough theory, but it is no more than a toy theory for our purposes. It is about time we turn to a theory with real world applications – a theory like QED. The classical Lagrangian for this theory is given by:

$$\mathcal{L}_{QED} = \bar{\psi}(i\gamma^\mu D_\mu - m)\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \quad (0.26)$$

with ψ representing electrons (as spinors), γ_μ -s being the so-called Dirac matrices which are necessary to represent spin-1/2 particles and $F_{\mu\nu}$ being the electromagnetic field.

Counterterms to the above Lagrangian are introduced in a completely analogous way to ϕ^4 theory. Recall that the electromagnetic field is also written in terms of A^μ , so that the modified Lagrangian comes out as:

$$\begin{aligned} \mathcal{L} = & -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + i\bar{\psi} \not{\partial}\psi - m_R\bar{\psi}\psi - e_R\bar{\psi} \not{A}\psi \\ & - \frac{1}{4}\delta_3 F_{\mu\nu}F^{\mu\nu} + i\delta_2\bar{\psi} \not{\partial}\psi - (\delta_m + \delta_2)m_R\bar{\psi}\psi - e_R\delta_1\bar{\psi} \not{A}\psi \quad (0.27) \end{aligned}$$

Now, the exact value they should take will come out from considering corrections to specific QED quantities. These can then be classified into three groups:

1. Photon Propagators: these corrections will involve loops interrupting the uninhibited photon “trajectory”⁵. From the interaction term of the Lagrangian we expect the loop

⁵Technically, there is no sense in talking about a spacetime trajectory for the photon in this case, at least in the “orthodox” / textbook understanding of quantum theory as no path has been recorded yet. For a more realistic framing of this, decoherence effects have yet to kick in.

diagram to have the following form (diagram adjusted from Schwartz 2014, p. 313): in




Figure 5: Photon Propagator Correction

which at any interaction point only two fermions and one photon meet. This process is called **vacuum polarisation** as the electron-positron pair in the loop acts as a virtual dipole in the vacuum. The corresponding integral is:

$$i\Pi^{\mu\nu} = -(-ie)^2 \int \frac{d^4k}{(2\pi)^4} \frac{i}{(p-k)^2 - m^2} \frac{i}{k^2 - m^2} \text{Tr}\{\gamma^\mu(\not{k} - \not{p} + m)\gamma^\nu(\not{k} + m)\} \quad (0.28)$$

Evidently, this is a tricky expression to compute but the main steps of the process are basically those specified in section (3.2). The extra complexity in QED (compared to scalar theory) is that we need to work with spinors and therefore we need to use the various trace formulas to compute the traces above. A detailed catalogue of the needed results can be found in Schwartz 2014, section 13.2 or Aitchison and Hey 2013, appendix J. After proper manipulations we obtain:

$$i\Pi_2^{\mu\nu} = i(-p^2 g^{\mu\nu} + p^\mu p^\nu) e^2 \Pi_2(p^2) \quad (0.29)$$

with $\Pi_2(p^2)$ containing all the divergence (regulator dependence). Using dimensional regularisation we can compute this:

$$i\Pi_2^{\mu\nu}(p^2) = \frac{ie^2}{6\pi^2\epsilon} (p^\mu p^\nu - g^{\mu\nu} p^2) + \text{finite} \quad (0.30)$$

where we have separated the part that leads to the divergence - notice the regulator ϵ in the denominator that we need to send to 0.

The correction in the propagator entails a corresponding correction to the Fourier transform of the Coulomb potential:

$$V(p) = e^2 \frac{1}{p^2} (1 - e^2 \Pi_2(p^2)) \quad (0.31)$$

To renormalise we want to impose a condition that will connect the coupling constant (here the charge e) to the Coulomb potential at a particular scale p_0 . This implies that the charge of the electron will change depending on the scale triggered. The condition chosen as our basis of comparison is:

$$\Pi(0) = 0 \quad (0.32)$$

which guarantees that the renormalised charge will be equal to the macroscopically estimated value. This will give a condition for counterterm δ_3 .

It is worth noting that these QED corrections to the Coulomb potential led to the discovery of the so-called Lamb shift phenomenon, which involves the splitting of the $2P_{1/2}$ and $2S_{1/2}$ levels of the hydrogen atom. The agreement with experimental data contributed to the acceptance of QED.

2. Electron Propagators: these corrections involve loop corrections to the line describing the propagation of a fermion (here: electron) and are described by the so-called electron self-energy diagram (diagram taken from Schwartz 2014 p. 322):



Figure 6: Electron Propagator Correction

Classical physics was long tantalised by the problem of self-interacting electrons (i.e. the electron interacting with its own field). Renormalisation appears to provide QED with the necessary arsenal to tackle this obstinate problem. The integral for the above diagram is:

$$\begin{aligned} iG_2(\not{p}) &= iG_0(\not{p})(i\Sigma_2(\not{p}))iG_0(\not{p}) \\ &= \frac{i}{\not{p} - m} (-ie)^2 \int \frac{d^4k}{(2\pi)^4} \gamma^\mu \frac{i(\not{k} + m)}{k^2 - m^2 + i\epsilon} \gamma^\nu \frac{-i}{(p - k)^2 + i\epsilon} \end{aligned} \quad (0.33)$$

in which the self-interaction graph Σ_2 using, for example, dimensional regularisation can be evaluated at:

$$\Sigma_2(\not{p}) = \frac{\alpha}{\pi} \left(\frac{\not{p} - 4m}{2\epsilon} \right) + \text{finite} \quad (0.34)$$

where all finite terms have been cobbled together to expose the real divergence. Note

that divergences here are of two types: one involving \not{p} and one m . So, two kinds of parameters will need to be adjusted in this case. According to the on-shell scheme the conditions imposed on Σ :

$$\Sigma_R(m_P) = m_R - m_P$$

$$\left. \frac{d}{d \not{p}} \Sigma_R(\not{p}) \right|_{\not{p}=m_P} = 0$$

guarantee that the renormalised mass will be equal to the pole mass or physical mass, and the residue equal to i . This helps fix the value of counterterms δ_m, δ_2 . Of course, if minimal subtraction were chosen, different conditions would have been applied.

3. Vertex: The third possible correction comes in the form of the vertex diagram, i.e. the diagram describing the interaction $e\bar{\psi}A^\mu\psi$ between electrons and photon (diagram taken from Schwartz 2014 p. 346): This diagram is denoted by $-ie_R\Gamma^\mu$. The calculations in




Figure 7: Vertex Correction

this case are even more involved so we will not quote any results here, but suffice to state the renormalisation condition is:

$$\Gamma^\mu(0) = \gamma^\mu \tag{0.35}$$

which essentially says that renormalised electric charge should be set equal to the value measured using Coulomb's law at sufficiently large (asymptotic) distances where radiative corrections do not come into play. This fixes the last counterterm δ_1 .

Renormalisability So, to recap, our process was the following: rewrite the Lagrangian with a bare and a counterterm, regularise the integrals, identify the infinite diagrams and then impose conditions to fix the counterterm part of the Lagrangian so as to eliminate the infinities.

Of course, eliminating the infinities in some diagrams is no guarantee that the theory will-behaved in full generality. How can we be certain that new complications will not arise as we attempt to calculate higher-order diagrams such as (diagram adjusted from Schwartz 2014 p.

384-5): Luckily, it turns out that for QED it is possible to give an argument that the process of




Figure 8: Higher-order Diagrams

parameter adjustment we went through for the basic infinite diagrams of the theory suffices to produce finite answers to all orders! Providing the whole proof is a rather involved task but it is possible to get some first estimate of whether a diagram will diverge or not using the concept of the **superficial degree of divergence**:

$$D = 4L - 2B_I \quad (0.36)$$

where L the number of loops and B_I the number of internal lines of the diagram. The idea roughly is that the divergence of any diagram will be exacerbated the more momenta one has in the numerator and will be alleviated the more momenta one has in the denominator. One then expects the following:

- $D > 0$: the diagram is divergent
- $D < 0$: the diagram is finite
- $D = 0$: the diagram will diverge logarithmically

There are various equivalent way that the formula for D can be written. For example, one implicating the number B of bosonic fields and F of fermionic fields, is:

$$D = 4 - \frac{3}{2}F - B \quad (0.37)$$

With this formula one can see, counting the relevant lines corresponding to bosons and fermions, that diagram (a) above has $D = -2$ and diagram (b) has $D = -1$ which are both lower than 0. Consequently, they are perfectly finite diagrams. The last diagram presented, (c) will turn out to be renormalisable through the recursive procedure we mentioned above.

Renormalisability is (or at least used to be) such an important property of a theory that theories were in fact classified based on whether it was possible to renormalise them or not.

Remarkably, it turns out that this fact can be read off a theory's coupling constants by checking their dimension, Δ_i . This is because the dimension of the parameter gives us a first rough estimate about the magnitude of the contribution of the diagram. We will see more on this in the following section when we examine the Wilsonian approach to the RG. In the end, we have the following possibilities for our theories:

1. Finite theory: it has no divergent diagrams
2. Super-renormalisable theories: it only has a finite number of superficially divergent Green's functions; its coupling constants all have $\Delta_i > 0$
3. Renormalisable theories: it only has a finite number of superficially divergent Green's functions; its coupling constants all have $\Delta_i \geq 0$
4. Non-renormalisable theories: it contains any infinite number of divergent Green's functions (as we go to higher-orders); at least one coupling constant has $\Delta_i < 0$

0.3.3 QCD

Having reviewed most of the facts of interest in the simpler cases of ϕ^4 theory and QED and given the complexity of introducing the formalism of QCD we will omit any detailed treatment of renormalisation in the theory of strong interactions. Most of the complexity comes out of the fact that QCD is a non-abelian gauge theory. Indeed, even the effective Lagrangian of a non-abelian Yang-Mills theory like QCD looks significantly more complex than the QED case:

$$\mathcal{L}_{QCD} = -\frac{1}{4}(F_{\mu\nu}^a)^2 - \frac{1}{2\xi}(\partial_\mu A_\mu^a)^2 + \partial_\mu \bar{c}^a (\delta^{ac} \partial_\mu + g f^{abc} A_\mu^b) c^c + \bar{\psi}_i (\delta_{ij} \not{\partial} + \not{A}^a T_{ij}^a - m \delta_{ij}) \psi_j \quad (0.38)$$

where many of the terms like F obtain additional indices because they are written in matrix form and additional ghost terms (called Fadeev-Popov ghosts) c, \bar{c} are introduced to help preserve the unitarity of the theory. One then writes down Feynman rules in diagrammatic form to describe the basic processes of the theory. An important bit is that we are no longer dealing with a theory of one mediator like the photon in QED but rather 8 distinct gluons, which can interact with one another. No wonder that QCD is fraught with indices!

Unsurprisingly, renormalising QCD is a rather tricky business requiring a lot of heavy technical work. There are many diagrams to be taken into account no less because gluons interact with themselves creating self-interacting loops like in the following graph (diagram adjusted from Schwartz 2014, p. 517):




Figure 9: Loop Diagrams in QCD

As a consequence, renormalising QCD even at 1 loop is a pretty challenging task (note that we also need to include a diagram (9b) for the Faddeev-Popov loops on top of the other fields!). It comes as no surprise that no less than 8 counterterms are needed to kill off all the infinities appearing at 1-loop. Yet, for all the complexity of this operation, the basic steps of the renormalisation procedure are the same as in the simpler cases we have already examined.

Further happy news awaits! QCD turns out to be a renormalisable theory much like QED (e.g. see Weinberg 1996, chapter 17 for the renormalisation of gauge theories). This means that one only needs to specify a set of parameters in order to derive finite results for divergent diagrams across all orders. Happily, all theories used in the construction of the standard model, that is, the electroweak theory as well, are renormalisable. In fact, renormalisability used to be seen as an obligatory feature any QFT should possess if it were to be an acceptable theory because non-renormalisability was associated with a failure of predictiveness. The argument was that if an infinite number of counterterms (each order putting more on the list) was required to tame the divergences in calculations, then there would be no hope of specifying all the necessary counterterms to produce finite results for non-renormalisable theories. Thus, whether a theory was renormalisable or not was elevated to the status of some selection principle. We will see that the advent of EFTs changed this prevailing mentality.

0.4 Renormalisation Group

At a first glance, various steps in the renormalisation procedure might appear arbitrary. And, perhaps somewhat surprisingly, this first impression is quite right. One can choose to regularise the theory using a cut-off, dimensional regularisation or whatever scheme they fancy and then eliminate the infinity along with a finite part they (also) get to choose. The specific choices do not matter as long as the physical results remain the same; what tools will be chosen is a matter of (usually computational) convenience. The renormalisation group is built upon this very idea: the choices one makes do not matter for the physics of the problem.

Now, there are two distinct ways one can think of the renormalisation group stemming from the two roles it played historically. The former, issuing from the work of people like Gell-Mann and Stueckelberg in the 1950s⁶ was more intimately connected to facilitating calculations in renormalisation theory while the latter, associated with the work of Wilson (K. G. Wilson and Kogut 1974, K. G. Wilson 1975), relied on stronger analogies with statistical physics to provide a deeper understanding for renormalisation as a physical procedure. We will examine each in turn and underline their particular uses:

0.4.1 Continuum Version

We start from the RG approach most favoured by physicists these days. The basis for the so-called continuum approach to the RG is the fact we have been repeatedly stressing up to this point that physical quantities such as correlation functions cannot depend on the choices we make when renormalising. In particular, the physics cannot depend on the particular scale μ we choose to define the renormalised quantities. This key insight is reflected onto the fact that parameters in a theory are not fixed but also change with the scale μ . In mathematical terms, if \mathcal{O} is the physical quantity of interest, this translates into:

$$\frac{d\mathcal{O}}{d\mu} = 0 \tag{0.39}$$

⁶For a history of the origins and development of renormalisation techniques starting with the problem of the stability of the electron and its self-interaction one can consult Brown 1994. For a history of the development of the two approaches to the renormalisation group to be discussed here see (J. D. Fraser 2021). It is worth quoting the following concluding remark by Fraser (p. 126):

The present study makes clear that the project of giving a more complete historical analysis of the renormalization group should not be conceived as tracing the development of a unitary scientific concept. Rather, an approach that is sensitive to the existence of many quasi-independent strands within the renormalization group tradition is needed.

As we shall see when we discuss naturalness in chapter 3, the “multifacetedness” of the renormalisation group will play an important role in evaluating the connection between naturalness and decoupling.

which simply states that no matter what scale μ , Λ etc one chooses to define their parameters, the physics (for instance, the relevant probabilities if we are talking about an amplitude) will stay the same. In fact, if we think of a state space where the axes correspond to the parameters of a theory (e.g. coupling constants λ_j , then a point will represent a particular set of values at a given scale and the RG will give a trajectory across a space of theories along which parameters change their values with scale.

The scale Λ , μ etc enters any renormalisation problem either implicitly or explicitly. When a hard cut-off like Λ is introduced, then scale enters in a fully explicit manner as the level at which one “cut off” higher energy degrees of freedom. In dimensional regularisation μ enters more subtly into the equations as a parameter introduced to ensure that the coupling constant of interest will be dimensionless in $4D$:

$$\lambda = \mu^{(d - \frac{d-2}{2})} \lambda_{4D} \quad (0.40)$$

Since the physics must remain unchanged, for S matrix elements evaluated at different scales μ , μ' , we must have:

$$S(\mu, \lambda_1, \lambda_2, \dots, \lambda_n) = S'(\mu', \lambda'_1, \lambda'_2, \dots, \lambda'_n) \quad (0.41)$$

In other words, all parameters can be adjusted so that the amplitude, the quantity we measure after all, is invariant under a change in the renormalisation condition (which is represented by the change in scale $\mu \rightarrow \mu'$). The field itself should also change under the scale transformation according to a relation of the form:

$$\phi' = \zeta \phi \quad (0.42)$$

Since fields at different scales are connected according to this relation, it is possible to establish the connection between more general, n-point functions at different scales as follows:

$$G^{(n)}(p_1, \dots, p_n; \lambda_1, \dots, \lambda_m) = \zeta_\phi^{-n} G'^{(n)}(p_1, \dots, p_n; \lambda'_1, \dots, \lambda'_m) \quad (0.43)$$

Now for the trick. We can take advantage of the scale that is always present in any renormalised theory to extract a very useful equation from the massless theory. Recall that since no scale appears in the bare theory, it must remain invariant as we change the scale. Of course, this means that if we take the derivative of any bare quantity with respect to any scale the result will

be 0:

$$\mu \frac{d}{d\mu} G^{(0)} = 0 \quad (0.44)$$

But now, recalling that $G^{(0)} = G^{(0)}(\lambda_1, \lambda_2, \dots, \lambda_n)$, i.e. the parameter dependence of the correlation function, we can write down an equation, known as the **Callan-Symanzik equation**, which relates the sum of partial derivatives of the parameters:

$$\left(\mu \frac{\partial}{\partial \mu} + \mu \frac{\partial \lambda}{\partial \mu} \frac{\partial}{\partial \mu} - \eta \mu \frac{\partial \eta}{\partial \mu} \right) G^{(n)} = 0 \quad (0.45)$$

where η is simply a rewriting of Z as $Z = 1 - \eta$. Typically one then defines the following quantities:

$$\beta = \mu \frac{\partial \lambda}{\partial \mu} \quad \gamma = -\mu \frac{\partial \eta}{\partial \mu} \quad (0.46)$$

where the $\beta(\mu)$ is the so-called beta function which describes the evolution of the parameters as a function of scale and γ is the so-called anomalous dimension which codifies the evolving scaling behaviour of the field. The C-S equation is then written more compactly as:

$$\left(\mu \frac{\partial}{\partial \mu} + \beta \frac{\partial}{\partial \mu} - \eta \gamma \right) G^{(n)} = 0 \quad (0.47)$$

The renormalisation group is a very powerful tool that allows us to examine the behaviour of theories across scales. One can input the corresponding parameters for theories like QED, QCD etc in the above equation to extract pretty valuable information about their value as μ slides. The β functions in particular can be studied to uncover the behaviour of theories at extremes:

- If $\beta(\mu) \rightarrow 0$ as $\mu \rightarrow 0$, the theory is **infrared free**, i.e. interactions become weaker and weaker at shorter distances; this is the case with QED, for example
- If $\beta(\mu) \rightarrow 0$ as $\mu \rightarrow \infty$, the theory is **asymptotically free**, i.e. interactions get weaker at larger distances; QCD is an example of such a theory
- If $\beta(\mu) = 0$ the theory is scale invariant and the corresponding point in the parameter space is called a **fixed point**.

These facts imply that perturbative techniques eventually breakdown for QED as we probe higher and higher energies (this associated with the existence of a so-called Landau pole in the

beta function of theory) whereas QCD faces the opposite constraint with perturbation breaking down at low energies (this is related to the confinement of quarks in hadrons). One of the main goals of the RG analysis is to track the evolution of a theory's parameters around fixed points. This is frequently presented in diagrams such as the following (diagram taken from Fradkin 2021, p. 418):




Figure 10: Renormalization Group Trajectories

Here, we have a theory with a parameter space consisting of h_1, h_2 . We see that in the first case (a) the parameter trajectories are converging towards the fixed point. This means that as we go to lower energies, i.e. as we approach the IR, the contribution of the processes represented by the corresponding operators will become increasingly negligible. By contrast, operators corresponding to parameters in the second case (b) flow away from the IR fixed point and thus will not become negligible at large distances⁷. As we shall see in the next section, operators are classified based on their "flowing" behaviour.

0.4.2 Wilsonian Version

The second approach to the RG, inspired by the work of Wilson, starts from a slightly different perspective. According to Wilson, we should imagine, as is reasonable to expect, that there is a scale Λ above which we are ignorant of the relevant physics; for all we know, the framework of QFT itself could break down. Instead of despairing, we try to turn this deficiency on our part to an advantage: we will simply disregard processes above the scale Λ imposing a hard cut-off to our theory. In the path integral formulation this can be represented as:

$$Z(\Lambda) = \int_{\Lambda} D\phi e^{-\int d^4x \mathcal{L}[\phi]} \quad (0.48)$$

⁷Of course, we could have reversed the flows such that we were tracking points of attraction as we move to the UV as opposed to the IR but this would not affect the analysis. Note that in case (c) we have operators of both kinds.

The idea now is to go over a coarse-graining procedure whereby higher energy degrees of freedom are “cast aside” and their effects incorporated into a redefined set of parameters. Crucially, the scale Λ here is treated as physical, i.e. a scale at which physics could become discrete (e.g. the string fundamental scale) and massively above the scales of interest. This is in sharp contrast with the scale μ in the continuum approach to the RG, which is always set close to the energy scales of the processes under investigation. Here is the Wilsonian recipe in more detail:

1. We introduce the physical cut-off Λ separating known from inaccessible physics like in $Z(\Lambda)$ above.
2. We set a lower scale $\frac{\Lambda}{b}$ inside the known physics and perform a split between high and low frequency modes for the fields: $\phi = \phi_h + \phi_l = \phi_{>\frac{\Lambda}{b}} + \phi_{<\frac{\Lambda}{b}}$. We want to “integrate out” the high frequency modes so:
3. We split the path integral into two parts:

$$Z(\Lambda) = \int_{\frac{\Lambda}{b}}^{\Lambda} D\phi_l e^{-\int d^4x \mathcal{L}[\phi_l]} \int_{\frac{\Lambda}{b}}^{\Lambda} D\phi_h e^{-\int d^4x \mathcal{L}[\phi_l, \phi_h]} \quad (0.49)$$

and then perform the integral only on the second (high frequency modes) part to derive a result that only depends on the low frequency fields:

$$Z(\Lambda) = \int_{\frac{\Lambda}{b}}^{\Lambda} D\phi_l e^{-\int d^4x (\mathcal{L}[\phi_l] + \delta\mathcal{L}[\phi_l])} \quad (0.50)$$

All dependence on high frequency modes has been transferred to the parameters of low energy physics.

4. We want to return to a theory of the same form as the original so we rescale the momenta $p' \rightarrow bp$ (note how this undoes the Λb rescaling) and then rescale the fields themselves $\phi(b) = b^{d-d_\phi} \phi'(p')$.

In general changing the cut-off Λ will lead to adjustments to the fields and the parameters. If we write our expansion as:

$$\mathcal{L}(\Lambda) = \frac{c_r(\Lambda)}{\Lambda^{4-r}} \mathcal{O}_r \quad (0.51)$$

we easily appreciate the fact that as Λ is taken to lower and lower values, the behaviour of the operators \mathcal{O} and their respective parameters c_r will change. Now, the dimension r plays a

crucial role here for it can let us create expectations for the way these will evolve. We can then classify the operators as follows:

- **Relevant** [for $r < 4$]: these are operators that will become more and more important, i.e. they will contribute more strongly as we reach lower energies
- **Irrelevant** [for $r > 4$]: these operators will tend to become inconsequential as we approach lower energies; they are said to be suppressed at great length scales
- **Marginal** [for $r = 4$]: the behaviour of these operators will vary and needs to be examined on a case-by-case basis

As an example, consider the following Lagrangian for a simple scalar theory. We write down the expansion containing all possible interactions with coefficients suppressed by appropriate orders of the scale involved (here Λ):

$$\mathcal{L} = \frac{1}{2}(\partial\phi)^2 + \lambda_2\phi^2 + \lambda_3\phi^3 + \lambda_4\phi^4 + \lambda_5\phi^5 + \lambda_6\phi^6 + \dots \quad (0.52)$$

Simple dimensional analysis on the terms comprising the Lagrangian allows us to see that the field ϕ is of energy dimension or $[\phi] = [E]$. This is derived from the kinetic term under the condition that the action which is the Lagrangian density integrated over spacetime $S \propto \int d^4x \mathcal{L}$ is dimensionless. Now, the dimension of the rest of the terms $c_k\phi^k$ is found based on these two facts (the dimension of ϕ and the dimensionless-ness of the action):

$$\begin{array}{lll} [\lambda_3] & \rightarrow & [E] \\ [\lambda_4] & \rightarrow & [1] \\ [\lambda_5] & \rightarrow & [E]^{-1} \\ [\lambda_6] & \rightarrow & [E]^{-2} \\ \vdots & \vdots & \vdots \\ [\lambda_k] & \rightarrow & [E]^{4-k} \\ \vdots & \vdots & \vdots \end{array}$$

We can see that the dimension of the coefficient for the ϕ^6 processes is of dimension $[E]^{-2}$. This means that this process will be suppressed by a factor of 2 in the cut-off. In other words, processes with coefficients of order $4 - k$ are not expected to play an important role at low enough energies since their effects will only be seen when energies are high enough for the $[E]^{4-k}$ suppression to be overcome. By contrast, processes such as ϕ^3 in this case which have a

[E] contribution are expected to dominate the physics at low enough energies. That is why the respective operators are called relevant.

0.4.3 Connection to Renormalisation

Conceptually, the most important takeaway from the Wilsonian approach to the RG is the way it illuminates the physical basis of the renormalisation procedure. Instead of viewing it as an arbitrary procedure for “sweeping infinities under the rug” (see J. D. Fraser 2020a for a discussion), the more physical picture drawn by Wilson allows us to see why some diagrams can be eliminated with no fear while other must be preserved. The key result is the characterisation of the operators in terms of their behaviour as the cut-off scale is changed.

The previous example naturally leads us to an appreciation of this fact. Once we have a Lagrangian expansion like 0.52 written down, we can see that higher-order processes are heavily suppressed by powers of Λ as we go to lower and lower energies. As a result, we can safely disregard their contribution to the (lower) scales of interest. The result? We can consistently work with the theory as if these effects are not there. But, this is exactly the part of the theory that is responsible for the infinities we came across when calculating Feynman diagrams. So, when we renormalise a theory, all we do, according to this approach, is forget about the intractable and unknown microscopic the theory cannot describe and restrict it to its appropriate regime (see also Burgess 2020, section 3.1.4 for a discussion).

If this re-conceptualisation of the renormalisation process can be retained, evidently it goes a long way towards alleviating any reservations about the ad hoc and purely formalistic nature of the elimination of infinities we saw in the previous sections. More on the philosophical importance of these issues in chapter 2.

0.5 Effective Field Theory

There is no single characterisation of an EFT - except perhaps for the fact that it is an “approximate” description of a certain level of physics in the (sometimes deliberate) ignorance of a more fundamental level of description. EFTs can be employed in a multiplicity of ways and providing a unified account is not entirely possible. Petrov and Blechman 2015 (sect. 1.5) offer a tripartite classification that relies mostly on the technicalities of the problems each type is meant to address. For example, the distinction between their Type II and Type III EFTs has to

do with how large the momentum transfers between interacting fields are. Here, we will offer the following, more conceptually oriented, classification that more adequately serves the goals of the rest of this thesis.

0.5.1 Top-Down EFTs

Perhaps the most intuitive way to understand an effective theory is in the top-down manner. Terminology here can easily become the source of confusion, so let's make clear what the “directionality” is meant to imply. The top theory is the more fundamental theory, i.e. the theory that penetrates into the finer details of smaller distances or higher energy scales. Producing a top-down EFT, then, is tantamount to “coarse-graining”: we start from fields and processes that are only manifest above a certain energy scale and replace them with an effective description in which we use lower energy degrees of freedom. It is as if we are blurring a picture. Mathematically, we split the Lagrangian of the more fundamental theory into two parts and retain the one corresponding to low energy fields:

$$\mathcal{L}_{full} \equiv \mathcal{L}_{\Lambda < + \Lambda >} \xrightarrow{\text{integrate out } \Lambda >} \mathcal{L}_{\Lambda <} \equiv \mathcal{L}_{eff}$$

Of course, higher-order physics cannot simply vanish out into thin air. The dependence of the effective Lagrangian on the higher energy physics is transferred or encoded in the (modified) coefficients or coupling constants of the new theory:

$$\mathcal{L} = \alpha \mathcal{L}_1 + \beta \mathcal{L}_2 + \dots \rightarrow \mathcal{L}' = \alpha' \mathcal{L}'_1 + \beta' \mathcal{L}'_2 + \dots \quad (0.53)$$

with $\mathcal{L}_i[\phi_1, \phi_2, \dots]$ “clusters” of operators like ϕ^4 , $e\bar{\psi}A\psi$ etc. The procedure by which this is achieved is called “integrating out” because we choose a fixed scale Λ and then we split each field into two modes:

$$\phi = \phi_{>\Lambda} + \phi_{<\Lambda} \quad (0.54)$$

and then perform the integral only over the higher energy modes of the given path integral much like in the Wilsonian approach to the RG. In this manner the higher energy degrees of freedom drop out of the new (effective) Lagrangian and do not appear as external lines in the processes. They, however, can still appear as higher-order corrections in internal parts of a diagram. This is a consequence of the uncertainty relations which allow for violations of energy conservation for sufficiently small time intervals.

As an illustrative example (see Petrov and Blechman 2015 4.1 for details) consider the following Lagrangian describing the interaction of light fermion field ψ of mass m and a heavy scalar field ϕ of mass M :

$$\mathcal{L} = i\bar{\psi} \not{\partial}\psi - m\bar{\psi}\psi + \frac{1}{2}(\partial_\mu\phi)^2 - \frac{1}{2}M^2\phi^2 - g\phi\bar{\psi}\psi \quad (0.55)$$

with g is the coupling constant determining the strength of their interaction. Implementing the aforementioned procedure for the the functional integral of the theory (assuming that we are probing energies sufficiently lower than the scalar field):

$$Z = N \int d\psi d\phi e^{i\int d^4x \mathcal{L}} \quad (0.56)$$

we can make ϕ drop out of the picture and write an effective Lagrangian at $\mu \ll M$ consisting only of the lighter field:

$$\mathcal{L}_{eff} = i\bar{\psi} \not{\partial}\psi - m\bar{\psi}\psi + \frac{g^2}{M^2}(\bar{\psi}\psi)^2 \quad (0.57)$$

Poor ψ can only appear as part of the correcting loops like the following for sufficiently small time intervals but cannot appear as an out-going (on-shell) particle at the low energies triggered (diagram taken from Petrov and Blechman 2015):




Figure 11: Heavy field ψ correction to light field ϕ propagation

To treat a more real world example, in Fermi theory, the low energy effective theory of the electroweak interaction, the suppressed high energy field will be the boson propagator. Indeed, at energies lower than about 100 GeV the underlying propagating W -boson field can essentially be replaced by a four-point interaction governed by the so-called Fermi coupling constant:

$$\mathcal{G} = \frac{\sqrt{2}}{8} \frac{g^2}{M_W^2 c^4} \quad (0.58)$$

Evidently, as long as the energies are pretty low the momentum of the field will be negligible compared to its mass $q^2 \ll M_W^2$ and the approximation $q^2 - M_W^2 \approx -M_W^2$ will hold (diagram taken from Burgess 2020).




Figure 12: From Electroweak Interaction to Fermi Theory

The fact that underlying processes are nicely suppressed as we “zoom out” to larger distances allowed physicists to successfully work with theories like Fermi’s up to some energy scale. The breakdown of the theory was frequently seen as a signal that new physics is around the corner that the currently possessed theory could not capture.

0.5.2 Bottom-Up

To make life easier, but, most importantly, to deal with cases in which the more fundamental theory is not known, another strategy of constructing EFTs has been devised. Clearly, if the more fundamental theory is not known, the above procedure cannot work: one cannot simply integrate out known degrees of freedom or include corrections of processes they are unaware of. The alternative strategy, which is known as **bottom-up** approach, starts out with writing down the most general Lagrangian possible for the specific problem (which means constraints) and empirically specifies the necessary parameters to make predictions:

$$\mathcal{L} = cK[\phi_1, \phi_2, \dots] + g_1V_1[\phi_1, \phi_2, \dots] + g_2V_2[\phi_1, \phi_2, \dots] + \dots \quad (0.59)$$

where K signifies the kinetic part and V_i the various interaction terms such as $q\bar{\psi}A^\mu\psi$. Empirically finding out what the parameters g_i requires a statistical analysis over measurements of processes involving the coupling constants. Unsurprisingly, this is far from an easy task (see Burgess and Moore 2006, appendix A for a brief discussion). One can then come up with estimates of the parameters within some desired margin of error.

At this point, and with good reason, one cannot help but feel that the whole procedure of writing down the most general Lagrangian and subsequently extracting any predictions what-

soever is a rather hopeless task. If anything, the problem seems intractable: keeping track of all the possible processes contributing, specifying parameters, estimating observables “in the dark”. There are a few things to consider, however.

First, the form of the Lagrangian is heavily constrained by symmetries or other principles that are imposed by the physics of the problem. So, if we assumed that a certain quantum number like the lepton number is conserved no terms violating this constraint will be included. Conversely, it might turn out that some terms must be included to ensure that meaningful results can be produced. This is the case with ghost fields included for gauge theories, for instance. Good insight into the physics of the problem plays a key role in simplifying the task at hand.

Furthermore, the “strength” of processes appearing in the expansion will depend on the energy scale. Power counting lets us know which terms can be ignored for a desired level of accuracy. Why bother calculating very complex processes that can only show up in energies that are not triggered by the experimental setup at hand? This strategy is a double-edged sword, however. On the one hand, it lets us simplify the computational task by allowing us to focus on part of a potentially huge Lagrangian. On the other hand, if mistakes are made in estimating what terms are relevant or not, the results obtained will be inaccurate. Such mistakes can be exploited to motivate progress in the field.

Finally, one does not always need to be ignorant of the more fundamental physics in order to favour a bottom-up EFT construction. Since the processes of smoothly going from the high energy to the low energy theory by integrating out the relevant fields might be too hard, an alternative is to construct an EFT by including the terms deemed relevant for the problem and then obtain estimates for the parameters by **matching** with the more fundamental theory. When matching two theories, one typically calculates some important quantity that both need to agree on such as some correlation function. Then, by comparing the two one can fix the relevant parameters of the EFT.

All in all the bottom-up approach can be summarised as follows:

1. Specify the most general Lagrangian consistent with the constraints (e.g. symmetries) of the problem. This can be an infinite sum.
2. Make sure to include all terms that are significant for the problem. Assess their importance (get an estimate about their contribution) using some power counting scheme.

3. There are two options now:
 - (a) Use experiments and statistics to obtain estimates of the parameters (such as coupling constants).
 - (b) Use a more fundamental known theory to compute some key physical quantities (e.g. a four-point interaction). Calculate the same in the aspiring EFT and force them to be equal at a specific scale. This will fix the values of the desired constants.
4. Use the new theory to make predictions about other quantities within the regime of applicability of the EFT. If things start to fall apart, then further processes might need to be taken into account.

There are plenty of examples of theories constructed using either option, but two of the most well-known are the Euler-Heisenberg, in which photons are seen as interacting at low energies because the more fundamental degrees of freedom (the propagating electrons between them) are not included and the Standard Model itself, the parameters for which are only experimentally specified. We will discuss both of these in some detail in sections 1.3, 2.5 respectively, so to spare the reader, we will refrain from repeating the same information here.

0.5.3 New Fields

There are cases for which one does not possess a way of “smoothly” patching together theories at distinct levels. It so happens that sometimes the more fundamental theory at a certain energy scale is known and we also possess a lower-scale effective theory but the degrees of freedom between the two are not the same. The theories are written in terms of different fields:

$$\mathcal{L}_{full} = \mathcal{L}_{full}[\phi_j(x); j = 1, \dots, n] \quad \mathcal{L}_{eff} = \mathcal{L}_{eff}[\psi_k(x); k = 1, \dots, m] \quad (0.60)$$

In this case one has to find a way to match between fields of a dissimilar kind in ignorance of the exact mechanism that makes the more fundamental fields give rise to the fields observed at larger distances. This usually involves a procedure whereby one imposes the symmetries of the more fundamental theory to the effective theory and invokes some symmetry breaking mechanism that results in bound states of the new degrees of freedom. This is the case with QCD, which is written in terms of quarks and gluons, and its lower energy limit that comes in terms of hadrons. We will say more about this in section 2.5.

Chapter 1

EFTs as “Theories”

EFTs have to occupy a central place in QFT. In fact, what we may call the “effective theories paradigm” has come to permeate all of contemporary physics, heavily influencing the physicists’ understanding of physical theories. Hence, it is no insignificant undertaking to investigate whether and to what extent existing philosophical accounts of scientific theories can appropriately describe their structure and account for their function in the broader edifice of physical science. Indeed, from a certain point of view the very term “theories” is a misnomer. For an EFT appears to be both richer and poorer in content than a standard theory. Poorer in that an EFT never incorporates the full relevant physics, i.e. it does not reach up to all energy scales, and thus never aspires to a Newtonian-style universalism. Richer too in that whole theories can be seen as EFTs within their restricted domain of applicability. In this sense, EFTs possess more content than a simple theory like EM: they contain information about the way physics changes across scales, the fields that become relevant at the various energy levels and even hints about the expected breakdown of a particular description. These features also imply that EFTs cannot be equated with the concrete descriptions of systems that are models; they contain more information than a specification of a system’s parameters would furnish. Along these lines, Hartmann 2001) suggests that EFTs occupy a novel, intermediate ground between theories and models. It will thus be an important task to situate EFTs in the space of “theories of theories”.

Conversely, it will be equally important to find out whether the widespread use of EFTs in physics engenders a need to modify traditional accounts of the structure of scientific theories. The novel aspects of EFTs might simply not be amenable to the kind of tools that philosophers have traditionally employed when discussing theories of physics (e.g. possible world seman-

tics). We will see that the two traditionally prominent schools of thought, the syntactic and semantic, fail to adequately encapsulate the richness of (perturbative) QFT. Insights from “less formal” or “pragmatic” approaches will offer the necessary supplements to a more complete understanding of the theory and the role of EFTs in it. In fact, we will see that turning to a more “dynamic” and “nuanced” approach to theories will allow for a deeper appreciation of the multifaceted character of QFT. In fact, this approach will prove to be more familiar than would seem at first glance. An interesting corollary of this discussion will be that the debate between Fraser (D. Fraser 2009, D. Fraser 2011) and Wallace (Wallace 2006, Wallace 2010) over the nature of QFT need not be construed as real disagreement at all. Instead, the two can be seen as emphasising two distinct goals within a wider program.

The plan of this chapter is the following. In section 1 we briefly review the main points of the two dominant formal approaches to the structure of theories: the semantic and the syntactic view. In lieu of a complete exposition of the cons and pros of either of these views, we will adopt a “plundering” stance. We will extract what is relevant for the view to be constructed in section 2 while highlighting their respective mismatches with QFT as practised by physicists. In the second part of this section we will attempt to extract some insights on theories from more “informal” approaches that emphasise the role of scientific practice in theory interpretation. The main lesson obtained from this discussion, namely that a less stringent relation holds between theory and experience, sets the ground for the hybrid view presented in section 2. There, theories are depicted as dynamic conceptual systems that come stratified into different levels of abstraction, essentially creating a conceptual hierarchy that progressively branches into more and more concrete structures at each level before connecting to reality. This branching involves coupling the theoretical framework at hand with i) additional peripheral assumptions (e.g. about the structure of space-time in cosmology, the form of the stress-energy tensor in GR), ii) other disciplines (e.g. combining relativity and thermodynamics, or relativity and mechanics, relativity and quantum information theory etc), iii) alternative mathematical tools (e.g. shifting from ODEs to PDEs, applying differential geometry tools to Newtonian gravity). The empirical content of the theory is to be retrieved from these more particularised sub-fields created through all this intertwining while insisting on the existence of a theoretical core they all share (the point that seems to be neglected in approaches like that of Cartwright 1983). Most crucially for our understanding of QFT, this conception allows, or better still welcomes, revisions even to the most theoretical parts of a theory, like its axioms, as a means of enhancing its

applicability or the accuracy and intelligibility of its key terms and principles. In section 4 this view is presented as a potential dissolution of the Fraser-Wallace debate, in which each side is understood to be pressing for complementary rather than strictly incompatible goals.

1.1 Standard Accounts of Scientific Theories

Here we briefly survey some of key aspects of the way scientific theories have been discussed in the philosophy of science literature and attempt to evaluate how QFT and EFTs fit into them¹. Starting with formal accounts, the syntactic and semantic, we claim that they fall short of accommodating QFT as is usually practised by physicists. When shifting to informal approaches, we welcome the way in which they highlight the main inadequacies of the formal accounts examined, but lament what we perceive as a lack of systematicity especially for theories in which mathematics plays a distinctive role.

1.1.1 Formal Approaches

Beyond the Syntactic View

The so-called received view of scientific theories has traditionally been the main representative of this approach. According to a well-known story, the Logical Empiricists were deeply influenced by Frege's and Russell's development of formal logical tools and sought to apply them to the study of scientific theories in an effort to clarify the empirical content of their terms and eliminate fruitless metaphysical controversies. The syntactic view of theories reflects the goals of this program. According to this approach, a theory is taken to be a set of propositions and corresponding terms that can essentially be broken down into two parts, i.e. we can essentially perform a bifurcation of the language into:

- a **theoretical part** which includes terms and propositions involving the various abstract posits of the theories like mass, fields, spinors, genes, instincts, beliefs etc whose meaning indirectly derives from experience: *“The mass of the electron is approximately equal to $9.11 * 10^{-31} kg$ ”*

¹For the rest of this chapter QFT will refer to the “mainstream” presentation of the theory prevalent in standard textbook presentations. Since an essential feature of this theory is perturbation techniques, it might have been more appropriate to call it “perturbative QFT”, but for the sake of brevity we shall omit this. Another underlying assumption will be that QFT in this sense finds its most mature expression within the EFT framework that emerged in the 1970s. In this sense, it is to be equated with what Wallace 2010 coins “conventional QFT”

- an **observational part** which includes those terms and propositions that are in a more direct connection to experience and can be accessed via some method of observation that, to a first approximation, must be deemed unproblematic: *“The pointer of the voltmeter has moved to 10V”*

First-order logic allows us to transcribe or re-describe the theory in a purely syntactic form by replacing theoretical terms with variables and appropriately reformulating sentences using quantifiers, logical connectives, constants etc to their formalised equivalent. To illustrate this, consider the following simple example from Healey 2017 (p. 124). Kepler’s first law states:

The path of every planet is an ellipse with the sun at one focus

This can be formalised using a procedure, known to all students that have practiced translations in first-order logic, like this:

$$(\exists!y, z)L(u, y) \& \forall(p)(\forall x) \left(P(p) \& L(p, x) \rightarrow (\exists n)[d(x, y) + d(x, z) = n] \right)$$

with x, y, z variables, $P(x)$ the predicate “ x is a planet”, u a constant (the sun), $L(x, y)$ the predicate: “ x is located at y ” and $d(x,y)$ standing for the “distance between x, y ”.

The goal is to reorganise or “reconstruct” the theory into a clear axiomatic structure that will reveal the logical relations between theoretical terms and will thus allow one to systematically manipulate the propositions in a formal manner, i.e. perform derivations using logical rules and the axioms defining a given field. The theory will thus be handled as a an abstract mathematical system with results (theorems, corollaries, models) not yet -at least- mapped to reality.

This is no satisfactory account of a a physical theory, however, as it is exactly this correspondence to reality that marks the difference between a physical and a purely mathematical theory. This critical connection is established with the help of:

- **correspondence rules** responsible for mapping the theoretical to the observational part - effectively endowing the former with its empirical meaning

This is exactly what is meant when we say that the theoretical terms have derivative meaning: these “bridge laws” are necessary to understand the exact role played by any theoretical terms introduced to the theory. An example of such a rule is:

mean molecular kinetic energy is temperature

which associates a theoretical term, mean kinetic energy of molecules, to an empirical term like temperature. Temperature can be measured using a device like a thermometer under specified circumstances. Note that the distinction between the theoretical and empirical part is sometimes arbitrary. So, while observational terms need to be either directly observable or easily measured through some kind of simple technique (Carnap 1974, p. 226-7), oftentimes theoretical terms might figure in the “empirical part” when interpreting a successor theory. The meaning of a molecule in classical mechanics and kinetic theory might be taken as established when investigating the empirical content of the term “wavefunction” or “condensate” in the quantum context.

Limitations The received view has been heavily criticised since at least the 1950s across all fronts: from the possibility of having a clear-cut distinction between theoretical and observational vocabulary (Kuhn 1962, Feyerabend 1962, Putnam 1962) to the over-reliance on first-order logic and the failure of axiomatisation in most (special) sciences (Van Fraassen 1980, Suppes 1977), the disregard for heuristic thinking and pragmatically driven idealisations or approximations found in applied mathematics (M. Wilson 2006, M. Wilson 2017, Wimsatt 2007) or the over-insistence on logical derivability - especially when deriving observational from theoretical statements (Stein 1995). When it comes to QFT the restrictiveness of the syntactic framework is evident in at least two ways: i) the emphasis on axiomatisation and ii) the intricacies of matching theory to reality.

First, the theory resists axiomatisation in a strict sense. One may only offer some broader principles that permeate the framework (like those presented by Weinberg 1995), but these fall short of systematically reproducing all results. The only successful axiomatic schemes that have been developed so far only cover non-interacting theories or simplified interacting theories in spaces of dimension lower than the four-dimensional physical spacetime. Contrary to non-relativistic quantum mechanics, we do not even possess a set of principles that would, if not capture its physical meaning, at least present us with a “QFT algorithm” to systematise the practitioner’s procedure². Worse still, the theory seems to be lacking in mathematical rigour

²This is the main complaint against the received or textbook formulation of non-relativistic quantum mechanics (e.g. Albert 1992, Maudlin 2019): we only possess instructions as to how to treat a system quantum mechanically but these instructions only deliver an impoverished physical understanding in light of the measurement problem. Contrast this with special relativity for which a pair of physical assumptions were sufficient to both explicate its novel physical meaning as well as act as the springboard from which all main results could be derived in a systematic way – as any introductory textbook on the subject shows. Contrary to QM, physical understanding and mathematical systematisation in the case of SR seem to be aligned. Yet, in QFT we even lack this much.

altogether. One of the main tools employed in the standard perturbative approach, the interaction picture, is simply ill-defined in the context of infinite degrees of freedom - the infamous theorem by Haag (Earman and D. Fraser 2006)! As if this wasn't enough, the very expansions we obtain through perturbation theory are not converging (Dyson 1952), but asymptotic series: truncating them leads to perfectly sensible results (after renormalization), but the series as a whole fails to sum to a finite result.

Apart from mathematical rigourousness, a proponent of the syntactic view will also have a hard time accommodating the various situational and opportunistic techniques with which the QFT framework is applied to specific physical problems. For example, in treatments of infra-red divergences in QED one adds a small mass to the photon to avoid problems with very small momenta. This "mass assumption" is only treated as a mere computational device and is set to zero as soon as the computation is completed. Perturbation theory itself is based on the assumption that incoming and outgoing states can be equated with free states at sufficiently long timescales even though this is not mathematically sound in light of Haag's theorem. In renormalisation one is presented with a lot of freedom in what regularisation and renormalisation schemes to employ. This choice usually made on the basis of convenience rather than principle. The same applies to gauge fixing. Given the physical equivalence of all the schemes this is of course no real problem. However, sometimes physical insights might be more salient in certain schemes rather than others. For example, in standard presentations of the Higgs mechanism (see e.g. Schwartz 2014, p. 576) no care is taken to present things in a scheme independent way: one works with a specific gauge choice (the unitary gauge) which makes the slogan of "gauge boson eats Goldstone boson" more explicit³. Further issues of consistency arise in the sums physicists work with. Even after renormalisation the perturbative expansion is shown to be a non-convergent series; yet, scientists do not refrain from exploiting their lower-order terms for predictions. Finally, taking EFTs into account further complicates things as now a lot more of a physicist's intuition is required to choose right degrees of freedom for a certain level of precision in the treatment of the problem. The matching procedure also implies that the kind and precision of predictions a theory can make will be heavily dependent on specific choices for the free parameters.

³Earman (Earman 2004) takes issue with this idea of a non gauge invariant field, mere "descriptive fluff" can be "eaten" by the gauge boson to produce its mass. Indeed, it is worrisome that prima facie the validity of the mechanism might depend on a particular choice of gauge

Beyond the Semantic too

The semantic view, the currently most popular approach to scientific theories, starts by criticising the emphasis on first-order logic and syntactic tools and, in broad terms, seeks to equate theories with collections of models. Depending on the way one attributes content to the term “models”, one obtains a different variant of the view. According to the **set-theoretic approach**, one starts with the arsenal of set theory and equips it with the right functions (and, possibly, relations) to represent key physical magnitudes as well as appropriate relations to represent the laws of the physical theory of interest. With this clear presentation one can then derive the possible models, which are then to be matched with concrete physical systems on the basis of an isomorphism (Van Fraassen 1980) or simply similarity (Giere 2004)⁴. An example along these lines is Suppes’ (Suppes 1957) formulation of particle mechanics. Note that despite the invocation of axioms and formal tools, the semantic approach is not restricted to first-order logic and only seeks to express rather than reconstruct the theory in the framework of set theory.

Another variant of the semantic view, the **state-space approach**, fixes models as possible trajectories on the phase space of a theory. The state space is an abstract mathematical space coordinated by the free parameters of the theory (like position, velocity or even charge, mass etc) with points corresponding to possible states of the system. The dynamical equations (such as Newton’s second law) define trajectories of possible evolution across the points of the phase space. A model in the context of classical mechanics, for example, will be a specification of the form $M = \langle \Omega, \mathcal{R}, s = (x_1, x_2, \dots, x_n), A_1, \dots, A_m \rangle$ with Ω some Euclidean space, R the field of values for the parameters, s the state, with x_i -s concrete values and A_j the possible physical magnitudes of the system. The connection with reality does not come directly, but needs to be established through a hierarchy of models: we start from very elementary models, akin to qualitative sketches of phenomena (how will a particle travel after a collision, for example). Next we need to make correspondences between these “sketches” and numerical models, which enable us to treat the problem mathematically (introduce relevant parameters like the angle of deflection, the incoming momentum etc). These phenomenological models will further need to be embedded into models of the theory (e.g. the quantities must obey laws of conservation, descriptions in terms of fields etc). Success or failure of these embeddings will determine its validity (French and D’Ecio Krause 2006, p. 68-74, Decio Krause and Arenhart 2016, p. 15).

⁴Giere’s more pragmatic approach to this connection squares with the less formal views we will discuss next. Various takes on the representational capacity of scientific theories are presented in Frigg 2020

Limitations For versions of the semantic view that emphasise axiomatisation, the limitations we highlighted for the syntactic view will still apply. In general the possibility of model construction in the way envisioned here will depend on the strictness of the concept of “model” employed. As (Miller 2016 has demonstrated for the state space variant, one will be hard-pressed to delineate possible models within the perturbative framework of QFT. We briefly summarise his argument here:

1. The empirical content of the theory is obtained by assigning a truth value to the physical statements of the theory via a satisfaction function⁵.
2. This is not possible when we do not possess an exact solution of the equations, for we cannot match a precise result to a measurement outcome.
3. This becomes a serious problem for the view when we realise that in many cases we do not possess any exact solution for the problem at hand and must rely exclusively on perturbation theory to obtain a result (e.g. three body problem).
4. The way out is to treat the perturbative solution as an approximation to an underlying exact solution to which the series converges. The convergence criterion guarantees that the satisfaction function will still apply.
5. Things become complicated in the case of asymptotic series that only approximate a solution in their first terms and then diverge as we continue adding terms; approximated function is not unique. This is often addressed with Borel re-summation, which ensures uniqueness of the approximated function given satisfaction of some conditions.
6. Alas, in QFT even that is asking too much! The technique fails in most important cases due to the presence of singularities (e.g. renormalons).

Miller’s suggestion is to modify our semantics so that they can tolerate imprecision in the theoretical value of a magnitude, i.e. we now demand that there is a theoretical error along with the measurement error. A model then would be an adequate representation when $(r - \epsilon, r + \epsilon) \subset (m - \delta, m + \delta)$ with r the theoretical estimate and m the measured value (ϵ and δ define the respective errors. As Miller acknowledges, this has the “counter-intuitive” consequence that “the empirical content of the theory simply has a limited, but rigorously established,

⁵This essentially asserts that the statement “the value of [given] physical magnitude Q is r ” is true, perhaps through observation or experiment.

precision” and that “the truth values vary with experimental precision” (p. 20). There will simply be no matter fact about whether a certain observable has a value at an order of precision exceeding that of the induced truncation error for the series.

Miller’s view thus seems to contradict that of J. D. Fraser 2016, namely that perturbative QFT furnishes approximations not models of reality. Fraser distinguishes his notion of approximation from idealised representation: in the former all one needs to do is obtain a function which describes a property of the system by taking a limiting relation from an existing model. This leads to an estimate of the property - physical quantity up to a degree of accuracy without necessarily constructing a consistent new model in the process, which is a prerequisite for an idealised representation⁶. QFT, on this account, approximates the values of quantities like correlation functions, scattering amplitudes, poles etc without presenting us with full models. Note the significant departure from classical mechanics: there is no “full” solution available that we perturb. As soon as we leave the highly idealised realm of free fields, we are inescapably forced to deal with perturbative sums.

Despite appearances, I think that both Miller and Fraser are both pushing for a similar conclusion: to encompass the full representational capacities of QFT our conception of model construction needs to be relaxed. Insistence on isomorphisms, classes of hierarchies and precise models heavily restricts the range of theories that the view can address. This is why we now swift shift to a brief discussion of some key elements of the less formal or pragmatic approaches.

1.1.2 Informal Approaches

The collective of approaches that, following Winther 2021 in his helpful taxonomy, we lump together under the banner of “informal” share a distrust towards what they perceive as the excessively formal character of the syntactic and semantic approaches. Finding the formal demands on scientific theories too stringent and taking the project of rational reconstruction⁷ as neglect-

⁶Fraser refers to an example by Norton that nicely illustrates the idea: the ratio of surface to volume of a sphere. Since the ratio is $\frac{3}{r}$ one gets an approximation of this quantity for $r \rightarrow \infty$ without having a model of a sphere with infinite radius

⁷It is perhaps worth noting here that the formal approaches examined were not done full justice – at least not in their own terms. Their goal is not describe scientific practice or deal with scientific theories in the wild. Instead, they attempt to “domesticate”, i.e. systematise and purify (in cleansing the conceptual apparatus of inconsistencies, lack of rigour or vagueness) the theory. The transformed highly idealised end-product will then be suitable for the extraction of philosophical lessons. Thanks to Niels Linnemann for his suggestion to expand on this.

ful of key aspects of scientific theories, the more “informal” camp shifts its attention towards actual scientific practice and the complex ways in which scientists make theories applicable to real world problems. The vastness of the field does not permit a thorough examination of the various views, so we will confine ourselves to the presentation of some key lessons relevant to our project.

Actual Practice One key underlying theme is to approach theories, or models, as they are used by scientists instead of reconstructing them along some ideal philosophical desiderata. This means that very often philosophers will need to confront competing formulations or varied contexts of application of the same framework (e.g. consider QM as practised by someone working in quantum information and someone working in quantum chemistry). The “informal” approach will be to treat all of them as instructive in understanding the content of the theory. Accordingly, they will mostly avoid favouring more idiosyncratic formulations of a theory that are more fitting to philosophical prejudices about, say, ontology or consistency like AQFT, Bohmian mechanics etc.⁸). If science appears to be pluralistic and disunified, the experimental interwoven with the theoretical and scientists often employ analogical reasoning (e.g. between optics and mechanics, statistical mechanics and QFT, or even gravity and fluid mechanics) as well as heuristics or principles of dubious status (think of Mach’s principle) then this is an indication that philosophers need to pay attention to all these informal aspects and appropriately contextualise their conclusions. As Wimsatt 2007 puts it (p. 27):

In our flight from monolithic and exceptionless logic of science we should not miss the many techniques that are wide but not universal in scope - “a toolbox of science”.

Evidently, there is a corresponding shift from the project of unity and grand theories to an ever-increasing focus on special sciences, situated reasoning and de-idealised descriptions. This is a direction the EFT framework most clearly has led physical research these days. As we shall see the construction of models along the EFT guidelines, which after all better serves our understanding of perturbation techniques, begets a philosophical account that will take into account all these informal aspects of the practice of constructing an effective description of a range of phenomena.

⁸Not that these examples are meant to somehow delegitimise work in these areas or question their validity. They are merely provided as examples of formulations that never enjoyed mainstream status within the community.

Relaxed criteria Consistent with the above is a certain relaxation of the criteria with which to judge the merits of a theory. For instance, recent works in philosophy of science (Frisch 2005, Vickers 2013) have emphasised the importance of inconsistent theories in scientific practice. This includes, according to both Frisch and Vickers, classical electromagnetism⁹, where application of the Lorentz force-law leads to particles leads to divergences, but also hybrid theories like Bohr’s model of the atom. In a similar vein, M. Wilson 2017 (chapter 2) emphasises the role of “descriptive opportunities” that scientists often employ to make a theory applicable to a series of phenomena like imposing macroscopic constraints to significantly reduce the space of possible solutions at a microscopic level. The key insight is that scientists often come up with theories in a “sub-optimal” form, which resembles a melange of algorithmic procedures or a heuristic patchwork. This was exactly what QFT looked like before the advent of EFTs, and to some extent still does (at the very least for proponents of axiomatic programs). Interpretation might be quite a toilsome task in this context, but the framework is constructed with efficiency as the primary goal. Disparate techniques might need to agree on some main results, but it might be impossible to derive a whole level of description from another, more fundamental level. We frequently need to rest satisfied with something akin to “coherence bonds” that render transitioning from one theoretical frame to another. EFTs greatly facilitate this task by delineating the relevant regimes of applicability and flagging the contributing and non-contributing degrees of freedom.

Models Philosophers in this tradition also speak of models (e.g. Cartwright 1983, Giere 1999, Giere 2004) as the basic units of analysis, but what they have in mind is different from the more formal concept we come across in the semantic tradition. Models are not taken to be (some formal) structures satisfying a theory, but are rather conceived of as basic representational devices, which scientists employ to “model”, i.e. describe and make predictions of the world. Models acquire a more “situated” or pragmatic status similar to that of maps (Giere 2006), whose accuracy is measured against the purpose they serve. This is why idealisations and approximations possess a central role: the model is tweaked to latch onto reality. Models are constructed by humans with limited capacities for perception and computation and thus are intended to maximise the output predictive power under these constraints (Wimsatt 2007, chapter 5). This is

⁹Naturally, this has sparked a big controversy among philosophers with Belot 2007 and North 2007 suggesting that Frisch’s argumentation is flawed. It is note that North finds Frisch’s “intriguing idea” as “dangerously close to accepting orthodox ‘Copenhagen’ quantum mechanics”.

interestingly similar to the way physicists have operated throughout the development of QFT and the standard model. The calculation of scatterings of electrons off protons was initially performed by treating the protons as structureless entities. One of the techniques employed today in lepton-hadron scattering is to treat the incoming lepton as a massless boson (Schwartz 2014, chapter 31). Weinberg himself highlights the importance of such more strategic, heuristic reasoning, in his discussing the role of renormalisability as a constraint for theory construction; after all, this allowed him to drop some (non-renormalisable) interacting terms and arrive at a simplified theory for the electroweak interaction. Even though the advent of EFTs and the rehabilitation of non-renormalisable theories, reduced the importance of renormalisability as criterion, its value as a guiding principle cannot be overstated. Neglecting it would adversely affect our understanding of the standard model.

Limitations For all the important parallels between the practice of QFT and the more informal approaches to scientific theories, there is a clear limitation that we need to recognise. These pragmatically-minded theories of theories” tend to be negligent of or even denying (e.g Cartwright 1983, Cartwright 1999) the systematic - unifying aspect of scientific theories. This, however, is a particularly glaring omission in the context of QFT, which was largely shaped by carefully imposing and revising symmetry (and other top-down) constraints. A key component, for instance, is the commonality between the way interactions are described in group theoretic terms. Abelian or non-abelian they endow us with a powerful recipe with which to construct Lagrangians: a part corresponding to propagating bosons, a part for propagating fermions and a part for the interactions of the two. Similarly, symmetry considerations led to the postulation of the right structures for quarks or the prediction of anti-particles. There was a fine interplay between the generalisation of the mathematical structures involved on the basis of experimental results and the creation of experimental expectations on the basis of mathematical principles. The edifice supporting QFT was largely raised on mathematical grounds. Forgetting about it means missing on a crucial aspect of QFT just as much as neglecting instances of “situated” reasoning. This is particularly true of EFTs which make essential use of symmetry principles to track the relevant terms in the effective Lagrangian and, of course, the RG transformations. In particular, with respect to the latter, philosophers (e.g. Morrison 2015) have even taken the universality property of RG techniques (which, remember, allows one to be ignorant of the exact nature of the fundamental theory) to be an instance of a mathematical explanation. At least

prima facie, a structural feature shared by some many disparate physical theories (from Ising models to field theory) is indicative of a deeper fact about the world.

1.2 A Plastic View

Combining the insights from the preceding discussion we want to suggest a hybrid, more dynamic view of theories that not only pays attention to all the situational factors contributing to the construction of models but also respects the importance of abstract theory building within mathematical physics. Here we focus on the way theoretical construction can proceed in a science like physics such that it will best enable contact with reality or phenomena¹⁰. It should be noted from the outset that the view developed here and the presented meta-model for scientific theories is assumed to capture aspects of scientific theorising within a broader range than the narrow confines of an EFT reading of QFT. In fact, one of the most important ideas we want to highlight in the following chapter is that far from constituting a dramatic break from the past, EFTs –or the EFT meta-paradigm– only bring to the fore aspects of theorising that have long been employed even in fields such as classical mechanics even though they were frequently ignored by philosophers¹¹. Still, EFTs add something novel to the mix: the technical apparatus necessary to implement some steps in a more systematic and quantitative manner. We will make these points more concrete in section 1.3, where the abstract discussion of this section will be mapped onto specific aspects of the EFT framework.

1.2.1 First Level: Stratification

Drawing on similar considerations found in Curiel 2019 it will be convenient to adopt the following schematic stratification of theoretical structure. Note that this is not a transcendental or normative schema of what a (physical) scientific theory is or should be - no claim to identifying essential features of a scientific discipline are presented. Rather, the approach is pragmatically-oriented: the descriptive account is vindicated to the extent that it most accurately represents key aspects of theorising and may be revised as needed to achieve this goal.

¹⁰Note that no commitment is made here to an unproblematic notion of observables or raw empirical data. Reality and phenomena might be as theory-laden as one desires, but this will present no problem here. Following Curiel 2019 we take it that experimental techniques and the experimentalists' more practice-oriented and situated insights play a crucial role in making sense of what a theory says about physical reality.

¹¹I am grateful to Marie Gueguen and Wayne Myrvold for pressing me to disambiguate this point.

1. **Frameworks** constitute the most abstract theoretical structures employed. They form the lens through which we “agree” to investigate phenomena. In a nutshell, they define a state space, rules for describing simple and complex systems, associations between mathematical structures and physical quantities as well as laws of dynamical evolution (see Wallace 2020 for a more detailed discussion). We can discern three main frameworks in modern physics are: a) the classical, b) the relativistic and c) the quantum. In the classical framework, all system descriptions will involve Newton’s law of motion and the Newtonian kinematic structure. The kinematics become Lorentzian in the relativistic context and the dynamical law is modified to transform appropriately under Lorentz transformations. The kinematical frame is significantly changed in quantum mechanics. In a quantum framework, physical quantities are represented by operators and commutation relations are imposed on them to codify their compatibility. Connections between the frameworks are sometimes taken to be determined on the basis of the behaviour of constants like the speed of light, taking $c \rightarrow \infty$ when going from relativistic to classical, and Planck’s constant, taking $\hbar \rightarrow 0$, when eliminating quantum effects¹².

2. **Theories** are formulated within frameworks to provide abstract descriptions or (high-level) representations of classes of phenomena such as electromagnetism (à la Maxwell or QED), gravity (Newtonian or relativistic), fluids and so on. They are basically the flesh we put on the bones of frameworks to obtain equations of motion that will describe the possible evolution of any system. So, in Maxwell’s (classical) theory of electromagnetism the four well-known equations describe the evolution of electric and magnetic fields and, along with Lorentz’s force law, can be used to track the evolution of a system in EM. Of course, in most realistic cases we need to combine different theories to produce a more accurate description of the evolution of a (realistic) physical system. Depending on the goals set for the inquiry, we might need to treat a mass as charged or as neutral, as a continuous distribution or a point particle, as a rigid or non-rigid body, take into account tidal forces and so on. It is the inter-blending of individual theories¹³ that enriches and truly discloses a particular theory’s empirical content. Thus, eventually one is led to “hybrid” disciplines like materials science, quantum optics, electronics or

¹²This is of course meant to be a caricature of the way the frameworks connect or reduce to one another. The quantum to classical transition is a vast subject and work in the area has made clear that the $\hbar \rightarrow 0$ or Ehrenfest theorem approximations are all but rules of thumb hiding a far more complicated task (e.g. Batterman 2001, Landsman 2017, Schlosshauer 2007). For interesting historical discussions about the relation between the classical and quantum frameworks from the pioneers of quantum theory see Bokulich 2008, chapter 2-4)

¹³At least if we grant textbook presentations as a good first approximation or road-map to such an individuating task.

cosmology, which truly reveal the explanatory and “enabling” character of theories.

3. **Models** comprise the most concrete representations of physical systems and can roughly be of two sorts: a) top-down and b) bottom-up. In the *top-down* case, one starts from an abstract theoretical description and tries to identify parameters relevant to a problem at hand. After an assignment of numerical values to these parameters it is possible to track the evolution of a system and make quantitative predictions. For example, one can assign a specific value to Hooke’s constant to obtain an equation of motion for a specific type of spring, specify the mass of two stellar bodies in order to derive their mutual attraction and orbits or even determine the free parameters in a QFT Lagrangian (consisting of the propagating and interacting terms of fields to obtain estimates of scattering amplitudes). Models vary in complexity depending on the representational accuracy one is interested in. Thus, when treating any particular system one often relies on simplifications that eliminate degrees of freedom that would heavily complicate calculations or incorporate them when necessary - e.g. neglecting the thermal properties of a wire might be ill-advised when constructing a delicate device like a computer.

In the *bottom-up* version, scientists might work without a specific theory guiding them explicitly, but might possess a lot of empirical data that they need to systematise. They might try to notice strong correlations between the various quantities implicated in an effort to extrapolate some deeper regularity at work. This might eventually lead to the development of a theoretical framework. An example of this would be the progression from Tycho’s observational data to Kepler’s “laws” and subsequently Newton’s theory of gravitation. It is here that scientists will often adopt heuristic and even inconsistent assumptions to make a first sense of a phenomenon (think of Bohr’s model or old quantum theory). The importance of these “bottom-up” models is that they are the first structured layer of data and thus the foundation on which to base comparison between theory and experience. Note how relaxed the notion of models has become in this context: we do not presuppose a coherent formal structure (phase space, axioms etc) to define them, but merely focus on their representational capacity and embeddability - manipulability within an abstract theoretical framework.

1.2.2 Second Level: Strategising

Theories should be seen in a more dynamic fashion. This can be understood as a certain open-endedness with respect to their more abstract components as well as their empirical content.

When it comes to the latter, we have to admit that only very rudimentary facts can be derived from a theoretical frame in isolation.

S1 Theories are almost always **patched** together, i.e. used in concert to describe a systems, and very frequently combined with pure empirical input, i.e. concrete empirical facts about features of the systems (not derivable theoretically) to do so accurately. The power of general relativity is only revealed when it is coupled with e.g. fluid mechanics to describe phenomena related to gamma rays and neutron stars - where relativistic effects need to be taken into account. Similarly, quantum mechanics plays a role in chemistry (e.g. with the description of molecular bonds) or information theory (with the introduction of qubit as the basic unit of information). As discussed by Curiel 2019 (section 4), the derivation a statement like “GR predicts the existence of a black hole at the center of the galaxy” will crucially depend on “...the confidence we have in the experiments and observations that delivered the data that allow us to identify and characterise SgrA* as a black hole, and more as *that* black hole” but also “...necessarily involves epistemic content from other theories - those e.g. we use to model the measuring instruments and the environment in which the experiment or observation is taking place”.

It is not hard to see that this procedure of patching theoretical frameworks depends on the particular goal and can assume multiple forms:

- **Parallel overlaps** between distinct theories such as when using a semi-classical approach for gravity in QFT, adopting fluid mechanics to Einstein’s theory of gravity or forming “hybrid” fields such as quantum statistical mechanics.
- **Vertical overlaps** between distinct levels of the same theory or theories grounded in the same frame such as Lagrangian treatments for friction and dissipative forces in complex mechanical systems.
- **Sub-model amalgamation** when a given phenomenon is compartmentalised, analysed using distinct frames and theories, which are subsequently woven together. Perhaps semi-classical techniques analysed by Bokulich 2008 can be viewed as an example.

In a lot of these cases there is no unified consistent framework that will allow a deductive-like treatment of the phenomena. Whether a unified framework can ever be found or will forever float in the skies of “in principle”-land is of course an open question, but our growing suspicions is that this has largely become an irrelevant issue. A natural consequence of all the

patching and overlapping between theories with theories and theories with empirical input is that there will hardly ever be a unique association of content to a particular theory. Like the meandering formation of a river and its tributaries, the ever-growing complexity of interwoven theories, practical techniques etc render the possibility of identifying a particular branch-like formation as a distinct theory a near practical impossibility.

S2 Deficiencies are not the only big challenge one may come across, however. Theories themselves will also often over-exceed their reach and make predictions in regimes that they were “not meant” to apply. This is the case with classical mechanics, of course, when it is treated as an all-encompassing theory. Its laws are simply inadequate to treat phenomena having to do with small distances and the stability of matter. Similarly, Maxwell’s theory fails when it comes to descriptions of atomic structure: the well-known Bremsstrahlung radiation will spell doom to any attempt at deriving a stable model of the atom. This is where a procedure we might call **tailoring** is required: any theory much like a roll of fabric needs to be cut and adjusted around the edges to properly match its intended domain. Allowing our theoretical fantasy to run wild is bound to lead to nonsensical results. This is what the various correspondence rules are supposed to do when it comes to the transition from the relativistic or quantum to the classical setting or when we need to abandon continuum mechanics for a statistical mechanics treatment of a system. They are signs – indicators of potential impasses, damaged asphalt, fallen bridges or unpaved roads.

In a somewhat loose sense, we can view this process of polishing our theories using patching and tailoring techniques as an operation almost analogous to the introduction of an affine connection on a manifold of theories. If each point represents a narrowly defined theory - although in agreement with the above the points should be thought of as blots in this space-, our strategising procedure enable us to make smooth transitions from one framework and theory to another, identify their respective domains (when new descriptive apparatuses become necessary) and compare results between the two. From a wider perspective these constitute techniques of “coherent-isation”.

S3 Previously, we alluded to a further open-endedness of theories with respect to their more abstract or formal (or mathematically-oriented) components. As we saw when examining informal approaches, it is a serious omission to be carried away by the pragmatic aspects of theories and neglect the role of systematisation brought by mathematical and formal tools. Apart from

considerations of an aesthetic kin (a more unified, all-encompassing framework certainly feels more appealing¹⁴) theory construction is also greatly facilitated when some basic, even not all-encompassing, principles are available to characterise and systematise a framework. One can think of the role of selection rules in quantum mechanics, Lorentz and gauge invariance along with renormalisability in QFT, the principle of equivalence in GR etc as relevant examples. Having a clear understanding of the physical and mathematical foundations of a certain frame is often a prerequisite for making important conceptual breakthroughs. As we will see, the cosmological constant problem itself could be the result of lack of rigour in the mathematical techniques utilised at the borderline between GR and QFT. Adopting a more flexible approach to systematisation along the lines of Stoeltzner 2004 we can recognise the value of what is called "**opportunistic axiomatisation**". Scientists often employ a set of provisional rules that will serve the purpose of endowing a theory with more cohesiveness even at the expense of rigour (think of Dirac's opportunistic expansion of the bra-ket formalism to position and momenta). If the right mathematical tools become available or significant physical obstacles are lifted¹⁵, then more satisfactory (i.e. broader and deeper) axiomatic systems will become available (again think of rigged Hilbert spaces and the GNS construction which put Dirac's heuristic tools on secure ground).

In analogy with the previous coherent-isation procedure, we may call this more formal strategising "consistent-isation". The goal is to harmonise the formal aspects of a theoretical framework and absorb any "internal shocks" produced by revisions that new empirical findings (such the discovery an anomalous new particle, deviations from expected scattering frequencies, violations of some conservation principle etc) induce.

1.2.3 Third Level: Combining

Let's pause and take stock for the moment. What is the view of theorising emerging from this discussion? In many respects the "plastic" view here is a hybrid view as well as a dynamic view. It is hybrid in that it constitutes a merging of formal and informal features. It is dynamic in that all parts of a theory are open to revision as response to challenges. Here is an allegory:

¹⁴Well, not so for everyone: Cartwright 1999 explicitly endorses an opposing view which she sees as more fruitful and minimising dogmatism in fields such as economics. However, I do not think that she would ultimately deny the validity of the more moderate proposal presented here.

¹⁵One could think of a revision of the QM postulates, for example. If a program like that of collapse theories is successfully completed, then we would naturally expect modifications to the dynamical equation or perhaps restrictions on the phase space.

...A group of brave settlers, who despise complacency and thirst for adventure, despair over the narrow confines of their community and seek out new opportunities for glory and expansion. After months of wandering in the open sea with vigilant eyes set on finding hints of new land, they finally spot a promising terra incognita in the horizon. Readily grasping the opportunity, they courageously settle the new land. At first, they need to explore, map and survive using whatever tools they have at their disposal. In the process, however, they refine their tools, become more well-adapted to their new home and begin to (re)distribute tasks among themselves to efficiently meet their new needs as they arise. Some become scouts and hunters: their goal is to bravely cut through the wilderness, explore new areas and furnish the community with food. Others prefer to cultivate the lands closer to home - a guaranteed way to avoid starvation [after all, wild animals sometimes escape their predators]. Others still, whose work often goes unnoticed or is sometimes underappreciated, stay at home to repair and construct tools, but, most crucially, design new, more efficient ones. Thus, life goes by in this thriving new colony with moments of frustration, disillusionment but also glorious days of success, abundance and celebration. Nevertheless, melancholic creatures by birth, our colonists eventually succumb to the nostalgic call of their old home at nights. Eventually, they dispatch envoys to their old community, which has also developed in the meantime, to share all about their venture. While a lot of what the envoys have to say sounds unintelligible or outlandish to their former comrades, the smartest among them can always resort to well-chosen analogues with local traditions to establish a common basis of communication. Regrettably, the hardest part is convincing them to start using the newly imparted knowledge and tools; old ways dies hard, as they say – sometimes they never do. Nonetheless, sufficient patience along with some tangible proof of their efficacy always -or always enough- suffice to make their case...

As with our adventurers so with scientists: some groups will be working on specific problems within a particular theoretical domain, trying to extract results to map onto data they have obtained. Others will try to combine techniques from different fields to more accurately describe a set of phenomena. Fewer will abandon a whole framework and redesign the whole edifice of physical research to deal with some “anomaly”. Scarcer still are these “under-laborers

of knowledge” who will try to put the techniques used by their colleagues on a systematic firm mathematical ground. Yet, all of these tasks reflect aspects of a scientific discipline and should be given appropriate attention if we are to understand the empirical or representational content of a particular discipline. We will now see how the effective approach to theories in physics may contribute to this analysis.

1.3 The Role of EFTs

Reviewing the relations between EFTs, theories and models, Hartmann 2001 claimed that EFTs occupy an intermediary ground in scientific practice and are not be equated with either of the other two. Theories usually aim at unifying explanations (subsuming a wide scope of phenomena under a set of more fundamental principles), constraining the possible models and providing global understanding for a particular system or phenomenon. By contrast, models aim at a furnishing local understanding, representing concrete features of a system and describing its evolution in time in causal terms. An example would be the relation between Newtonian gravitation and particular models of the solar system or QCD and the parton model for a description of proton - electron interactions. For Hartmann, EFTs partake in both of these theoretical constructs but fall short of the universality desired by theories (after all they are explicitly constructed to include only some degrees of freedom) but are also less constrained than models in that they are presented as widely encompassing descriptions up to a certain scale (think of Fermi theory or even the standard model).

While this approach to EFTs is a good approximation to an account of their importance and function in contemporary physics, it is also somewhat restrictive in that it relies on a somewhat artificially sharper contrast between theoretical structures, the limits of which are fuzzier than first thought. In particular, as Hartmann himself admits EFTs are sometimes hard to distinguish from models (his example being current algebra and Chiral perturbation theory), while, we might add, theories themselves are not so distinct a category in an era of physics in which even the most fundamental of theories (like QFT and GR) are to be treated as effective. Within the more dynamic picture of theories that we sketched above, however, a certain conception of EFTs can highlight the continuous nature of the suggested stratification and reveal the natural progression from frameworks to models.

The suggestion is to treat EFTs as part of a **meta-framework**, i.e. a certain conception of

what it means to do physics and in particular physics within the confines of the QFT framework. Its foundation is the idea that nature can be studied in multiple levels of description in which one only needs to take into account the degrees of freedom that are manifest or dominate processes at a given scale. This is an expectation formed by standard practice in science: we assume that details about quarks will be irrelevant in the domain of molecular biology much like fine details about the neurochemistry of the brain will play a limited role in the study of economic decision-making. Most of the time in physics we seek to do away with detailed descriptions of systems e.g. by treating bodies as point particles, forgetting about the underlying structure of neutrons or avoid tracking the micro-interactions of the basic constituents of stars. The effective theories meta-framework is a formal expression of this natural expectation.

Here is a sketch of how this framework squares with the more plastic picture of theories developed previously:

1. To construct an effective theory, we need to adopt a specific framework. In the context of EFTs in particle physics this framework is QFT. This places clear constraints on the kind of Lagrangians (essentially: theories) one can write down through principles such as Lorentz invariance, gauge invariance or renormalizability. For example, mass terms of the form $m\psi^*\psi$ may not be included for fermion fields as that would violate the $SU(2)$ gauge invariance of the theory. This is one of the reasons we need the Higgs mechanism to define the masses of the various fermion fields in the standard model.
2. A theory is now defined by a particular Lagrangian one writes down incorporating specific fields and interaction terms. To combine theories, that is, to take into account the various ways in which a system can interact with others, one needs to add appropriate terms in the Lagrangian consistent with the principles dictated by the framework. For example, one can treat the electromagnetic fields as massive by adding a corresponding mass term:

$$\mathcal{L}_{massEM} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \frac{1}{2}m^2 A^{\mu 2} \quad (1.1)$$

or one can ignore the electromagnetic interactions of quarks by including only terms relevant to how they interact under the strong interaction. One may also modify the standard model Lagrangian to include mass terms for neutrinos with the addition of an appropriate number of parameters. In each case, one is free to modify the theory they

are working with¹⁶. This is a more precise and “quantitative” form of “patching” that we discussed above. In this framework combining theories amounts to carefully choosing your interaction terms consistent with symmetry principles¹⁷.

3. An EFT can also be used to simplify calculations in accordance with the level of precision required. Often, one will construct a Lagrangian that does not contain all the possible processes involved but rather truncate it so that it contains only a certain number of them up to a particular scale of interest, with rest being suppressed by the energy scale. For example, an EFT for a scalar fields in this case would like:

$$\mathcal{L}_{eff} = \partial_\mu \phi \partial^\mu \phi + m^2 \phi^2 + A\phi^4 + B\phi^6 + \dots \quad (1.2)$$

with ... representing terms that have been omitted. These terms will be of the form $\frac{C_i}{\Lambda^n}$ with Λ the energy scale. The coefficients A, B will need to be fixed so that the theory produces accurate predictions. This is often done through a process called “matching”.

4. This matching process, in which the coefficients in the effective Lagrangian are adjusted to render the EFT an accurate description of interactions up to a scale, reveals the delicate interplay between theoretical and empirical - experimental aspects of a theory. This is because, broadly speaking, there are two ways to perform the “matching”:

- (a) The “theoretical way” is to compare the EFT to the full theory in order to extract the coefficients needed for making predictions. Consider the case of QED and its effective theory, the Euler-Heisenberg theory, which are respectively described by the Lagrangians:

$$\mathcal{L}_{QED} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \bar{\psi}D\psi \quad (1.3)$$

$$\mathcal{L}_{E-H} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \frac{\alpha^2}{m_e^4} \left[c_1(F^{\mu\nu}F_{\mu\nu})^2 + c_2 F^{\mu\nu} \epsilon^{\mu\nu\alpha\beta} F_{\alpha\beta} \right] + \mathcal{O}(m_e^{-6}) \quad (1.4)$$


where contributions from energies at order higher than m_e^4 have been suppressed. To “match” the two theories, one needs to calculate enough observables in the full theory to obtain a good estimate for the unspecified coefficients in the effective

¹⁶While we are here treating the narrow case of field theory, it is possible to imagine the extension of this framework to different areas of physics that can be given a Lagrangian formulation. Changes to the Lagrangian will correspond to different theories with which one attempts to describe the evolution of systems.

¹⁷Evidently, this does not exhaust the content of patching as we developed it above, which has a more pragmatic and multi-framework nature.

theory. Here there are only two such coefficients: c_1, c_2 . We start by computing the following diagrams in both theories (diagram taken from Becher 2015): Note how

QED computation



Effective theory




Figure 1.1: Euler-Heisenberg Effective Theory

the propagators for the more energetic electrons collapse onto points in the effective theory. Comparing the results obtained by the two theories we conclude that setting the coefficients equal to:

$$c_1 = \frac{1}{90} \qquad B = \frac{7}{90}$$

will make the effective theory consistent with the full theory up to this energy scale. We may now use this theory to compute more complex observables without being dragged into the complexities of the full theory. The only caveat is that we may not exceed the energy scale at which our matching procedure was defined.

- (b) The more empirically inclined way to construct a (simplified) EFT is to estimate the coefficients through a comparison with experimental results. In this case, no comparison between observables in the full and effective theory will be necessary. Instead, one estimates the coefficient through some scattering experiment (in a manner similar to what researchers do for when fixing the standard model parameters). Using a process such as trap experiments or various particle decays it has been possible to set highly precise bounds to the possible values parameters such as neutrino or pion masses can take. In the case of coupling constants like that of the electromagnetic interaction comparisons are made between the magnetic moment of the electron in experiments for the Josephson effect or the quantum Hall effect and

the theoretical estimates by QED (Burgess and Moore 2006, p. 482-4). Clearly, when full theories are unavailable or computations in them intractable this interplay between theory and experiment becomes an indispensable tool to extracting meaningful results from a theory.

Of course, none of the steps in the construction of an EFT constitute a straightforward enterprise: what fields to use, what effects to expect, what terms to neglect all depend on the kind of description that will balance computational tractability against explanatory completeness. Clearly, a good deal of “physical intuition” will be necessary for this task. Physicists frequently come to realise that they have neglected significantly contributing terms, which alter the predictions the theory makes. For example, when considering a low energy effective theory of QED with $E_\gamma \ll m_e$, we might neglect terms like:

$$\mathcal{L}^{(6)} = \frac{c}{m_e^2} j^\mu j_\mu \quad (1.5)$$

which will not contribute in the interactions of low energy photons $\gamma\gamma$. By contrast, this term will be important if we wish to consider corrections to the Casimir energy and omitting it will lead to incorrect results.

5. A novel feature of the effective theory meta-framework in QFT is the way in which it can systematically track the behaviour of theories as energy scales change. With the renormalization group we can set theories as part of a broader space coordinated by the free parameters in the Lagrangian. With the values of these parameters changing as a function of scale, the theories themselves will occupy various points of this space of theories indicating their behaviour in the UV or IR. Thus, equipped with the RG techniques we have a tool to patch the various theories with one another in a systematic way. We know, for example, that QCD has a Landau pole at around 100 MeV, therefore any approximations at low energies will fail to produce accurate results. Similarly, QED formally can only be trusted up to its own Landau pole at 10^{286} eV. However, we also know that QED can only be seen as an effective approximation to a more encompassing theory covering the combined standard model interactions and gravity, whose effects will eventually become important way before the Landau pole of QED becomes relevant.

Takeaway Having reviewed these key features of EFTs, we can now come to a concrete appreciation of the admittedly abstract discussion of the preceding section. We now see how viewing theories through the lens of the effective meta-framework can reveal their dynamic character. EFTs enable us to tell a more refined story about the way an abstract formal frame will connect to reality by establishing a continuum of progressively more encompassing descriptions. Starting from the abstract frameworks that determine the most general aspects of a theory like the kind of mathematical structures we employ, we form concrete theories through the specification of appropriate Lagrangians. We see how we can essentially mix different theories by incorporating terms that couple different kinds of fields with one another (like the fermionic field of electrically charged particles to the photon or the Higgs field). At the same time one pays attention to the effects that are significant for a given scale e.g. by turning propagating bosons as point interactions in accordance with violations of the uncertainty principle. By tracking the breakdown scales of the (now understood as) effective theories, we obtain a better sense of how we can patch them to one another.

1.4 What is QFT?

In this section we draw on our previous discussion to ponder about the nature of QFT as a theory. As we are about to see, viewed through the above lens, QFT as a discipline is richer and more accommodating than what philosophers have sometimes taken it to be. In fact, the two main trends represented by mathematical physicists and particle physicists can be said to implement different strategising procedures: opportunistic axiomatisation and patching-tailoring respectively. Far from rendering QFT a singular case, our treatment of theories underlines its commonalities with other subdisciplines in physics that are typically deemed unproblematic.

Dissatisfaction with the apparent lack of physical insight and the ad-hoc character of the procedure led to a reconsideration of the QFT project. One branch of research focused on the foundations of the theory in an effort to produce a rigorous mathematical treatment of fields that would eliminate the pathological divergences and render the theory consistent up to arbitrary scales. Frequently dubbed “axiomatic” QFT, the main idea behind this project is to present a set of principles or axioms similar to that of NRQM which will be sufficient to derive physically meaningful results without recourse to mathematically ill-defined tools. As an example, the Haag-Ruelle theory is a framework that describes scattering in a rigorous

manner eschewing the typical heuristic way of identifying Heisenberg fields with the incoming and outgoing fields at $t \rightarrow -\infty$ and $t \rightarrow \infty$ respectively. The Haag-Ruelle theory can be derived by a set of axioms proposed by Wightmann describing the state space of the theory, the field content and how it relates to the particle content (see Duncan 2012, 9.2-9.3 for details).

The **axiomatic approach**¹⁸ has so far garnered minimal support from the physics community with only a small group of people systematically working on it. The reasons for this appear to be two-fold. On the one hand, physicists often adopt a more opportunistic stance with respect to the theories and tools they utilise. They often sidestep worries about the mathematical consistency or the status of their approximations by emphasising the empirical success they enjoy. This, of course, has the advantage of maintaining a steady pace of progress without getting bogged down in the intricacies of mathematical “sophistication”. On the other hand, the emergence of renormalisation group techniques and the effective approach to field theories for the most part assuaged (or at very least were taken to assuage) the worries of the past. At long last, there was a way to understand renormalisation that was physically motivated and “rationalised” the dubious techniques of infinities subtraction.

The unpopularity of axiomatic approaches in the physics community was not shared by philosophers, however. Much work in the philosophers’ corner was based on axiomatic QFT, even though its minimal appeal to physicists was often recognised (Halvorson and Mueger 2006). The reasons for this voluntary neglect of practiced QFT were its perceived deficiency in what were considered to be essential desiderata for an interpretable theory. Mainstream QFT evidently lacked rigour, a clear mathematical foundation and a set of principles characterising it:

There remains an implicit working assumption among many philosophers that studying the foundations of a theory requires that the theory has a mathematical description. (The philosopher’s working assumption is certainly satisfied in the case of statistical mechanics, special and general relativity, and non-relativistic quantum mechanics.) In any case, whether or not having a mathematical description is mandatory, having such a description greatly facilitates our ability to draw inferences securely and efficiently. (Ibid)

¹⁸A word on terminology: by axiomatic here we mean any approach to QFT that seeks to cure the theory of its infinities by redefining it in a mathematical rigorous way. Thus, we are not committed to any program in particular be it constructive QFT or algebraic QFT or other. Further, although dubbed “axiomatic”, as D. Fraser 2011 puts, it echoing the voices of its founding fathers, this approach is not committed to axioms “set in stone”, but rather acknowledges the need for revisions given relevant empirical input. From a philosophical perspective, this is exactly what we would expect from a more “formal” way of thinking about theories that would prioritise clear mathematical formulations to be given a subsequent interpretation.

Despite its “demise” one can clearly perceive the, not so subtle, echoes of the received view here: a theory starts or should start or ideally is meant to start with a clear mathematical framework that philosophers can then go about interpreting with the help of formal tools be they first-order logic, model theory or, more recently, category theory. It is no surprise, then, that philosophers have for the most part shied away from the task of interpreting QFT as is standardly practiced by physicists. The tide has turned in recent years, however, as works have spawned in this direction taking seriously the idea that QFT should be viewed through the lens of the effective program. Instead of being a bug, renormalisation techniques are taken to be the key to revealing something deeper about physical theories and their connection to reality.

1.4.1 The Fraser - Wallace Debate

The contrast between the standard or mainstream approach to QFT, let’s call it LQFT (following Wallace 2006), and that of the axiomatic framework, which we shall abbreviate as AQFT (“A” standing for axiomatic rather than algebraic), has been the focus of an interesting debate between Wallace and Fraser. The question each of the two tries to answer is: which framework, AQFT or LQFT, should be the focus of philosophers’ attention? Wallace 2006 began a defence of “naivete” emphasising the vast empirical success the standard non-rigorous LQFT program has achieved in particle physics – with its extremely accurate descriptions from electromagnetic interactions to the constitutions of nucleons and beyond. He took this success and the apparent stagnation of axiomatic approaches as a clear indication that philosophers should shift their attention away from AQFT to problems involving LQFT:

...the problem with restricting our foundational studies to AQFT is that – pending the discovery of a realistic interacting AQFT – we have only limited reason to trust that our results apply to the actual world, which appears to be described rather well by the Standard Model.

D. Fraser 2009 picked up the gauntlet and undertook the task of making the case in favour of the philosophers’ preference towards AQFT more explicit. She did so by presenting the choice between AQFT and LQFT¹⁹ as an underdetermination problem: what we are presented with are essentially two rival programs which, starting from different standpoints about what a (good) theory is (or better: should be), adopt different strategies on how to deal with the

¹⁹Actually, Fraser’s choice in the 2009 article presented a trilemma between the axiomatic approach, the cut-off approach and an additional infinite approach to QFT. For our purposes, and the debate with Wallace, the first two will be sufficient.

infinities arising in QFT calculations. However, between the two, only the more rigorous AQFT approach, Fraser claims, can truly be said to encapsulate the spirit, true goals and scope of what QFT truly is because only AQFT fulfills two essential desiderata:

1. **QFT = SR + QM:** AQFT is the program that truly fulfils a central goal of QFT since it is the only framework that truly unites special relativity with quantum mechanics; only AQFT tells us what a world, both quantum and relativistic, would be like.
2. **Consistency:** AQFT is also the program addressing rather than avoiding the inconsistency of the interaction picture; it does not invoke non-rigorous or conceptually dubious tools; everything is perfectly well-defined

Fraser's criticism were addressed by Wallace in a more recent paper of his Wallace 2012 in which he assumed a more aggressive stance towards the AQFT project, essentially moving away from the more moderate coda with which he concluded his Wallace 2006 paper:

From this viewpoint, we can see that Lagrangian QFT (as I have defended it) is not really in conflict with AQFT at all. Success in the AQFT program would leave us with a field theory exactly defined on all scales, and such a theory would be a perfectly valid choice for 'theory X': furthermore, even if we found such an exact QFT it would not prevent us from defining low-energy, 'effective' QFTs – which would not be well defined without a cutoff; nor, probably, would it obviate the need for these theories in describing certain low-energy limits of X.

towards the following more bellicose intrata:

But forty-five years have passed. And they have seen theoretical physicists (notably Kenneth Wilson and John Kogut, in the early 1970s) approach the problem of renormalization from a very different direction. In doing so, these physicists made assumptions which directly contradicted some of the basic assumptions of the AQFT program: crucially, they assumed that instead of being definable on arbitrarily small spacetime regions, quantum field theory would break down at some short lengthscales. [...] Given these more recent developments, it is no longer appropriate — if ever it was — to see AQFT as the proposed mathematically-rigorous version of CQFT. The two are better understood as rival research programs, trying in different ways to resolve the problem of renormalization.

The central piece of evidence that Wallace credits for this change of mind is the deep conceptual breakthrough of the 1970s by people like Wilson, Kogut and others who developed the RG

techniques which eventually led to the prevalence of EFTs all over high energy physics. The goal of the rest of this chapter is to examine the force and implications of this claim. For the sake of completeness, but most importantly to motivate and facilitate the upcoming discussion, we summarise the key argument made by Fraser and Wallace in their 2011 exchange in the following table:

Table 1.1: QFT: AQFT or EFT?

Fraser	Wallace
AQFT seeks to provide a consistent set of principles that fulfil the goal of uniting relativity and quantum mechanics, i.e. QFT = SR + QM.	We should not expect theories to hold up to all scales (main lesson from EFTs). LQFT has been tremendously successful in explaining phenomena within its domain of applicability.
AQFT is not plagued by mathematically dubious techniques like renormalization, on which LQFT has relied upon to produce results.	Any dubiousness or ad-hocness in the renormalisation procedure was addressed by the work of Wilson and others.
Renormalization group techniques only serve to shed light on the empirical content of the theory. They leave us in the dark about its theoretical content.	AQFT wrongly expects that QFT will be valid up to arbitrary scales. We know this to be false given QG, final theory or even tower of EFTs.
It is unclear how to interpret the cut-off in LQFT. It is mostly used as a formal tool to aid in calculations. No clear physical lessons to be drawn yet. AQFT provides a clear interpretation along these lines.	The AQFT program has failed to deliver any results for actual theories. It has only succeeded in a very narrow range of theories in 2D or non-interacting theories in 4D.

There are no no-go theorems for AQFT. Why not focus on a program that has already been so successful? All we argue about is the relevant merits of two distinct programs. Overall, its virtues establish AQFT as the more attractive alternative.

There are many fruitful ideas promoted by its side, but let us specifically highlight three overarching themes of greater interest for our project:

1. **The nature of theories** Fraser's argumentation assumes that theories need to be cohesive frameworks that, at least in principle, entail "clear-cut", fully specified models of reality: there must a straightforward or almost deductive connection between axioms, first principles and derivable models. This contrasts with a more situated conception of theories that tends to view them as systematising schemes of a more approximate kin. This is reminiscent of the broader contrast between the formal and informal approaches: should one focus on conceptual and mathematical clarity and precision or on scientific practice and the way it establishes connection to the world? D. Fraser 2009 herself nicely captures this contrast in interpretive exigences (p. 558):

By "interpretation" I mean the activity of giving an answer to the following hypothetical question: "If QFT were true, what would reality be like?" In contrast, the interpretive question that Wallace focuses on is "Given that QFT is approximately true, what is reality (approximately) like?"

For Fraser we only possess a theory when we can meaningfully ask what the world looks like under the assumption that the given conceptual system is true. Clearly, LQFT fails in this respect because it explicitly denies addressing this question up to arbitrary scales. To illustrate how demanding this criterion is consider whether such a question can be meaningfully asked even in the context of classical physics. Does, for instance, classical EM or mechanics work up to all scales? For instance, classical models of the electron or the atom were clear failures in this regard leading to infinities (e.g. when trying to estimate the self-energy of an electron) or incorrect predictions (unstable electron trajectories).

Perhaps all we need in this case, according to Fraser, is that this description is possible, i.e. that the theory furnishes a description even if this is completely wrong. LQFT would be different in this sense in that it is inherently restricted due to renormalisation.

2. **Actual vs Potential Merits** This can be probably be subsumed under the broader question of what a scientific theory is. More specifically, when do we have a theoretical whole that is ripe for interpretation? Wallace claims that developments in 70s have rendered LQFT a framework mature enough to demand our full or, most of our, attention. The AQFT proponents insist on a certain level of mathematical maturity before we can get on board with interpreting a theory as used by physicists. Clearly, the outcome of this side of the debate will depend on the flexibility of one's standards or meta-theoretical claims on theories.
3. **The Nature of the Cut-Off** Should we think of the cut-off as a mere mathematical tool or as a physically meaningful element of the theory? This is a major issue that will impact the way we will think of the ontology of QFT. We will treat it separately in the following section.

To get a clearer sense of the stakes of the debate, it will be helpful to distinguish between two questions:

- Q1 Which program should be equated with QFT? Is the completion of AQFT a mandatory step before we can claim to truly possess a relativistic version of quantum mechanics? Or is the current less rigorous framework physicists work with adequate, perhaps even more informative than any complete axiomatic system?
- Q2 Which program should be the basis for our interpretation of QFT? Should we treat LQFT only as an instrument to perform calculations waiting for the rehabilitated AQFT version of the theory? Or should we start with whatever framework physicists use as our primary (or, perhaps, only) focus?

Reply to 1: I think that the basic lesson from the the discussion in this chapter is that, rather than seen as rivals, AQFT and LQFT are complementary projects within the broadly construed research program of QFT. Within a spectrum of possible theoretical undertakings, AQFT is closer to the most theoretical end of the spectrum and LQFT the more practice-oriented end.

The AQFT project is about constructing a consistent conceptual and mathematical framework, marrying special relativity to quantum mechanics in a coherent manner, while LQFT is interested in cutting through what is seen as “mathematical weeds” to derive concrete results. The latter naturally comes at the price of occasional sloppiness, while the former is frequently intercepted by swarms of formal intricacy. There are both advantages and possible drawbacks of this line of thought:

Table 1.2: QFT: Both AQFT and EFT

Advantages	Disadvantages
<p>It agrees with scientific practice. People working in AQFT do not see themselves as rivalling mainstream QFT, but rather putting it on firmer ground (e.g. exchange between Gross and Jaffe on proving that QCD is mathematically sound in Cao 1999, p. 164-5)).</p>	<p>What about developments that might render the axioms adopted obsolete or weak to capture the whole framework? What of the existence of competing axiomatic schemes?</p>
<p>It allows for our a more open-minded approach to theories, one that, in turn, allows drawing on the merits of distinct traditions - research directions simultaneously. Similar to concerns about the dominance of string theory in the post standard model approaches, there is a case to be made for the value of pursuing alternative projects: one that furnishes mathematical sophistication and the other explanatory power.</p>	<p>What about the conflicting physical meanings of the two programs? Surely, they cannot both be acceptable when they disagree on how one should think about high scale physics, the nature of the cut-off and the domain of applicability of QFT.</p>

Further research will reveal whether AQFT is a viable program if and when it is completed or if and when no-go theorems are proven. Either way this will greatly influence the way we think about aspects of QFT like the vacuum, the cut-off, interacting fields.

What happens when we develop a successor theory to QFT? In what sense will it be compatible with the AQFT program? The effective reading naturally leads to the idea of a more fundamental theory.

When it comes to the first objection, I think that it rests on a very stringent notion of what axiomatisation should be. As we saw, the proponent of AQFT will typically be open to a revision of the principles adopted and adjust them according to need. This is in line with the “opportunistic axiomatisation” we discussed a couple of sections ago. Just as the “practitioner” typically uses an abundance of heuristics to apply a theory to real world systems, so can the theoretician afford necessary amendments to render their axiomatic system strong enough to encompass as many aspects of the theory as possible.

The second and third objections will be treated as part of the last section on the physical meaning of cut-offs.

Reply to 2: Given our answer to question 1, it comes as no surprise to say that philosophers should not feel restricted to explore the implications of any one particular framework. Results from both axiomatic as well as practiced QFT have brought important insights in their respective domains. Duncan 2012, for example, is a nice example of a conceptual treatment of QFT that blends results from both traditions. Nevertheless, if, like in sports, watching a match end on a draw is somewhat unfulfilling, we may wish to declare the 2006 Wallace the winner of this debate. Philosophers should start paying more attention to the QFT as is practiced by the vast majority of scientists and in particular to problems that emerge in the EFT program. Such problems will have to do with naturalness, the transition to BSM physics, the applicability of EFTs, effective approaches to gravity and so on. Reasons not to neglect the AQFT program, however, include important results such as: rigorous derivations of the spin-statistics theorem (with the axiomatic structure allowing a better articulation of the assumptions needed for the derivation),

a series of no-go theorems that have implications on the status of particles in QFT (Malament 1996, Reeh and Schlieder 1961), the importance of inequivalent representations (e.g. Ruetsche 2011).

1.4.2 The Status of the Cut-Off

The success of the AQFT program and the correct physical interpretation of the LQFT approach essentially hinge upon how one understands the role of the cut-offs imposed in renormalisation procedures. Per objection 2, it seems that the physical lessons to be drawn are conflicting: if the AQFT program succeeds, then QFT should be consistent up to arbitrary scales and LQFT is wrong in treating QFT as only restricted to some energy scale Λ . On physical grounds, the two should be treated as competing programs. Yet, the persuasiveness of this argument depends on the way it is fleshed out. In particular, we need to understand the true significance of treating QFT as valid only up to a certain energy scale. It might turn out that this carries less ontological implications than one might *prima facie* anticipate. The key claims here are:

1. No clear lesson follows from the introduction of cut-offs for the ontological import of QFT. Spacetime might be discrete or continuous. Continuous RG does not eliminate any degrees of freedom; it just tracks their contribution.
2. The AQFT framework can also incorporate renormalization and the RG as a formal tool for easier calculations. This does not necessarily retract from its significance in the LQFT case. Constructive QFT still retains its physical significance while striving for rigour.
3. When a successor theory emerges, both LQFT and AQFT will be treated as approximations to it. They might differ in how this is spelled out. LQFT will necessitate a patching procedure: seeing where to delimit its regime of applicability. AQFT will look like a superseded framework: appropriate correspondence relations will be required. The physical significance will be the same, however: the world appears to be made up of fields only up to some level of coarse-graining.

First of all, it is not clear that the success of renormalisation techniques in LQFT implies a picture of reality that is incompatible with AQFT. That would be the case, for example, if the cut-off were to be taken as an indication that spacetime is not the continuous structure presupposed in AQFT, but rather discrete - like a lattice. As D. Fraser 2009 (p. 552) stresses it would

be very peculiar for QFT to be dictating something interesting about the nature of spacetime. One would accept theories of spacetime like quantum gravity to be our best consultants in this domain. And, indeed, such an interpretation is not forced on us. It relies on a literal reading of the Wilsonian approach to the RG, which more closely mimics the procedure of integrating-out small distance degrees of freedom in statistical mechanics. While a lot can be learned through comparison of the two, it is not self-evident that we may rely on statistical mechanics for any conclusions in QFT²⁰. In the continuum RG, however, which is also the version most often and conveniently used by physicists, there are no cut-offs. Rather, the point of the procedure is to see how the various terms in the Lagrangian expansion contribute as a function of energy scale considered. In this case, higher-order effects are simply shown to become irrelevant at low energies instead of being explicitly integrated out. No clear hint of a energy-momentum cut-off that would, in turn, imply a corresponding discretisation of spacetime follows.

“But”, the detractor persists: “what sense can we make of the RG in an AQFT context? In the former we are explicitly talking about relevance at separate energy scales and the whole axiomatic program is to make the theory consistent up to all scales. This is clearly as sharp a contrast as it can be!” Well, not necessarily. D. Fraser 2011 (p. 131) remarks that RG techniques only reveal the empirical rather than the theoretical content of QFT. In this sense, the cut-off should be treated as a formal tool, a part of a recipe with no physical significance. If the algebraic program turns out to be a successful project, one will likely be inclined to treat the cut-off as mathematical tool that is only meant to facilitate calculations. However, there are multiple programs that seek to inject mathematical rigour in QFT that do not jettison or fully revise the Lagrangian project. Programs within constructive QFT, for example, start from a version of Lagrangian QFT and try to put it on a firm mathematical basis by defining the measure for the path integral or connecting the perturbative expansion with some nonperturbative construction through techniques like Borel summability (Hancox-Li 2017, section 4). Algebraic QFT could thus be seen as the most general top-down program of “rigourisation” which abstracts from particular models to “derive structural relations among elements of the theory” (Fredenhagen, Rehren, and Seiler 2007, p. 64). The goal is to construct a consistent framework for QFT from first principles.

By contrast, the goal of constructive QFTs is clearly less ambitious and more “local”. The

²⁰The role of analogical reasoning in the context of QFT and statistical mechanics is of particular interest. A lot of the physical insights from the latter are transferred to the former theory given their shared mathematical structure. This topic certainly deserves closer analysis, but we will postpone this discussion for the moment. See section 3.5. where the task will be undertaken along with considerations of naturalness.

nature of the project is bottom-up: start from a Lagrangian like that of the ϕ^4 theory and construct an axiomatic framework, which will be satisfied by the theory's interacting models. It turns out that in such more "intermediary approaches" Fraser's claim that RG techniques only function as a calculation tool is too dismissive. As Hancox-Li 2017 stresses (p. 354):

It is not that constructive QFT on its own already has theoretical content without Lagrangian QFT, but to decide even what its theoretical content is, namely, its solutions, we need Lagrangian QFT for guidance. In other words, even if one insists on interpreting the contribution of Lagrangian QFT as merely "empirical," it is still the case that the empirical content is guiding what counts as acceptable theoretical content.

Li 2015 also examines the role of RG techniques in constructive field theories to claim that, even in this context, they offer information for the behaviour of the theory at the UV limit. He takes this as evidence that the RG should not be interpreted in formally or in a purely instrumentalist manner:

...axiomatic QFT is at best a kind of partial characterization of the theoretical content of QFT. Indeed, mathematical physicists have long acknowledged that CQFT provides additional dynamical information that a pure axiomatic approach does not.

It might turn out that all we can accomplish in the project of "rigourisation" of QFT is some variant of the constructive approach rather than the more abstract goals of the algebraic framework. The more narrow scope of this approach perfectly chimes with what we treated under "opportunistic axiomatisation": the project of a formal systematisation can be decomposed into small such sub-projects which might or might not coalesce into one coherent super-framework.

Before closing this section, let us add a few remarks about the relation of AQFT and a successor theory. First, even though EFTs have been taken as a natural standpoint from which to speculate on future physics, one cannot rule out the possibility that a future more fundamental theory will simply render the RG arsenal obsolete. Perhaps we will experience a fundamental shift in perspective with the introduction of a whole new framework - much like what happened with quantum mechanics back in the early 20th century. We might also witness something analogous to the introduction of fields in the 19th century or the abandonment of a fully mechanistic model of nature. It is also possible that, instead of being fully superseded, AQFT might bear a relation to the future theory X similar to the way Dirac conceived of the relationship between classical and quantum mechanics (see Bokulich 2008, chapter 3 for an account): that is distinct

frameworks that constantly inform one another to become progressively more encompassing and refined.

1.5 Quick Application: Scientific Representation

There is one further interesting consequence of treating theories as effective descriptions and this has to do with their function as representational devices. In particular, one of the standard problems in the philosophical literature on scientific representation has been the grounding of the distinction between accurate and inaccurate representations (Frigg and Nguyen 2020, p.12-4). Tolstoy's account of Napoleon in "War and Peace" is a largely inaccurate representation of the French emperor; one had better read some historical biography to form a well-grounded opinion. One cannot make a proper assessment on wrong data. Inaccuracy, however, is not always a bad thing. Treating the a planet as a point particle might greatly facilitate calculations about its trajectory whereas trying to fully represent its shape, volume, atmosphere, chemical composition etc would only render the problem intractable.

The problem essentially comes down to the following question: "what to base this omission of details on"? The pragmatist reply is to claim that theories are predictive, descriptive or systematising tools employed by (cognitive) agents with specific goals. Much like a map, a theory will only include so much detail and so many features of the target system as they will be enough for the desired goal. Accordingly, abstractions and idealisations are not only inevitable but also desirable features of this process of simplification and specialisation of the investigative problem. One complaint or worry about this more pragmatic mindset is that it runs the risk of making the whole procedure somewhat arbitrary: anything can be a representation, even a completely accurate one, as long as the model serves to satisfy some agreed goal.

There is definitely some truth in this goal-dependence of models, but we now see a way to make the whole process of abstracting from the details of a system more "objective". The framework of EFTs allow us to track in a quantitative manner the degrees of freedom that are relevant to a specific problem, the strength of their contribution, their effect on the behaviour of the theory, the regime of its applicability etc. Once a goal is set, that is, once we decide on the energy scale and level of precision we are interested in, facts about which degrees of freedom contribute and to what degree become determinate. Once the power counting begins, we can put quantitative estimates on the kind of processes that will be relevant at the energy scale of

interest; the representational capacity of our model is fixed independent of our declared goals.

Chapter 2

EFTs as Theories of the World

Scientific theories are meant to be more than mathematical *jeux d'esprit*. Their ultimate goal and significance is to describe, represent or, to avoid any more loaded term, “say” something about the world. In this chapter we turn to a discussion of the ontological implications of the EFT program. The motivating question for our inquiry is: what should we take the effective status of our theories to be implying about the underlying structure of reality? Somewhat differently, what is the most warranted, in light of accuracy and scope, picture of the world that emerges when QFT is approached through the lens of the EFT program? Are its implications for ontology as radical and dramatic as envisioned by Cao and Schweber 1993, for example, or should we heed the more sceptical or cautious voices -like those supportive of the AQFT program- about restricting the physical significance we attribute to it?

Before we are able to provide an answer to this question, however, it is necessary to revisit the all time classic debate on scientific realism. The goal is not to settle this dispute for good, but rather to investigate the ways in which this new meta-stance towards scientific theories might lead to revisions of the standard realist thesis. This issue naturally arises given the more restricted scope and inevitable lack of knowledge associated with effective descriptions. Accordingly, adopting a realist mindset within this frame seems a highly non-trivial task. Uncovering the representational capacities of theories will inevitable involve facts about scale dependence, computational complexity restrictions and more broadly the regimes of applicability of a particular theory. We will see that this will eventually lead to an abandonment of what we will call a fundamentalist attitude to theory interpretation: roughly, the idea that theories need to be construed in a purely bottom-up way specifying some fundamental constituents from which to reconstruct the “manifest” image.

We thus conceive of the project of this chapter as two-pronged. First, we will engage with recent work in the philosophy of EFTs (Fraser. 2018, J. D. Fraser 2020b, Williams 2019a), which seeks to motivate and defend a version of “effective realism”. This view is a essentially more restricted version of the realist thesis taking into account the unique features of effective theories. We will try to assess to what extent this is a viable position and whether it can be shielded against attempts at “empirical appropriation” (Ruetsche 2018, Ruetsche 2020a). We will see that a promising realist line, consonant with the rejection of fundamentalism, will be close to a view inspired by Stein 1989. The second aspect of the project involves an exploration on the possible ontological readings compatible with the EFT program in increasing order of ambitiousness: from instrumentalism to a denial of reductionism. The final verdict on the correct approach to take will have to wait subsequent investigations into the limitations of the EFT framework with respect to the hierarchy problem and the cosmological constant problem.

The plan of the chapter is the following. Section 1 begins with a historical discussion of the role of dissatisfaction with renormalisation procedures as the springboard for the development and acceptance of RG methods. This is taken as motivation for a Steinian form of realism. Section 2 takes on the issue of effective realism: it motivates, defines and refines the notion of realism in light of the EFT program emphasising its prospects for an exegesis of the notion of approximate truth. Two main challenges for ER are flagged: i) its incapacity to fully address the issue of theory change and ii) its possible collapse into some form of empiricism. The first challenge needs to be accepted as a unavoidable restriction of the realist thesis while the second is avoided through a rejection of fundamentalism. Section 3 takes up the task of clarifying this claim. Several possible ways of construing “fundamentalism” are presented and then assessed in light of the EFT program. Section 4, then, proceeds to redefine realism in a way that is not too demanding on the metaphysical side while remaining committed to the realist idea that theories reveal facts about the world.

2.1 Realism Matters

There is often a tendency within the scientific community to downplay the importance of certain philosophical debates. The debate over scientific realism will probably strike the unsuspected as somewhat unnecessary (in that the right answer is obvious), pedantic (all detailed discussions of the status of abduction, the definition of observables etc will most likely try the patience of

someone eager to jump to calculations and concrete data) or simply irrelevant (the “why care” attitude: in the end it will not make a difference to the way scientists do science). Scientists will often take on a pragmatic attitude towards their theories eschewing questions related to their interpretation (think of quantum mechanics), the metaphysical underpinnings of terms used (think of the quantum state, quantum fields) or the alleged equivalence of distinct formulations of the same theory (e.g. the choice between Lagrangian and Hamiltonian mechanics).

This pragmatic attitude will of course sound disappointingly instrumentalist to any philosopher with realistic sympathies. However, physicists have frequently, even “subconsciously”, adopted a more ambivalent attitude that oscillates between realism and instrumentalism. The historical development of renormalisation and the scientific community’s evolving attitude towards it attests to this. For, despite the immensely successful application of the various renormalisation techniques we reviewed earlier, a sentiment of dissatisfaction continued to pervade the physics community over its true status and deeper physical meaning. The main complaints raised against this whole procedure was its apparent arbitrariness underscored by the lack of a physically motivated rationale to support it. Up until the emergence of an EFT understanding of QFT around the 1970s, there seemed to be no clear answer to what J. D. Fraser 2020a has coined the “justification problem”: what is about these “odd” techniques involving ad hoc cancellations of infinities that allows them to produce some of the most accurate results ever produced in the history of physics. Adopting a more realist lens, the riddle can be alternatively put: what is it about the world that renders renormalisation such a successful tool? What do we learn about the structure of the world once we attempt to understand the physics underlying this procedure?

To motivate this question, consider the example of perturbation theory as it is employed in standard treatments of astronomical phenomena within the framework of classical mechanics. When trying to estimate the trajectory of a body like Jupiter orbiting the sun, one quickly notices that the complexity of the problem, which involves taking into account the gravitational pull of the sun as well as the other planets or even asteroids acting on the body of interest, will thwart any attempt to derive an analytic solution. Accordingly, one resorts to perturbation theory in the hopes of at least obtaining an approximate answer. The process is well-recorded (e.g. Goldstein, Poole, and Safko 2002, chapter 12). One first splits the Hamiltonian of the problem in two parts: one covering the tractable aspect of the problem -amounting to a solvable problem like the two-body problem involving Jupiter and the Sun- while relegating the rest of

the effects to a small, weakly contributing, correction part. Then, one expands around the solution of the first part treating the rest as perturbations that slightly modify with the degree of precision depending on the level of computational complexity we want to tolerate.

The success of perturbation theory in this case can be physically justified or explained by the fact that gravitational attraction falls off with increasing distance or decreasing mass, effectively rendering effects from bodies progressively less relevant the less massive or further distant they are. It is, then, expected that when investigating the interaction of Jupiter and the sun we can neglect the effects of distant stars, nebulae and galaxies, despite their massiveness due to their distance while small planets like Mercury or even smaller formations asteroids can be treated as minute corrections to the main solution. It is exactly this kind of physical insight that is lacking in renormalisation. Why neglecting the effects of processes produced at energies above a certain scale (with the cut-off) or simply dropping infinities (with dimensional regularisation and the MS scheme) leaves behind a finite quantity that just happens to be the one we find empirically? The mystery deepens when one realises there is a lot of arbitrariness in the choices made in all steps of the renormalisation process: choosing the regulator, the order of correction and the renormalization condition.

This fact did not escape the attention of many prominent physicists involved in the creation of QED. Even someone like Feynman, famous for his “shut up and calculate” attitude with respect to philosophically-inspired problematic, appeared acutely aware of and deeply dissatisfied with the misgivings of renormalisation as understood up until 1970s. In fact, even as late as 1982 he presented the problem in pretty negative terms as this quote from his popular book on QED suggests (Feynman 2014, p. 128-9):

The shell game that we play is technically called ‘renormalization’. But no matter how clever the word, it is still what I would call a dippy process! Having to resort to such hocus-pocus has prevented us from proving that the theory of quantum electrodynamics is mathematically self-consistent. It’s surprising that the theory still hasn’t been proved self-consistent one way or the other by now; I suspect that renormalization is not mathematically legitimate.

It is unsurprising, then, that someone like Dirac always attracted to mathematical elegance and arguably excessively driven by aesthetic criteria in theories would have similar feelings to express (quoted in Kragh 1990, p. 184):

I must say that I am very dissatisfied with the situation because this so-called ‘good theory’

[quantum electrodynamics] does involve neglecting infinities which appear in its equations, ignoring them in an arbitrary way. This is just not sensible mathematics. Sensible mathematics involves disregarding a quantity when it is small – not neglecting it just because it is infinitely great and you do not want it!

We see that both physicists are deeply dissatisfied with what they saw as an illegitimate or, at best, ill-defined mathematical procedure. For how can one make mathematical sense of a cancellation of an infinite quantity by another infinite quantity, which, if anything, seems to sweep infinities “under the rug” to produce a finite result? Dirac rightly commented that such operations are only legitimate for infinitesimal quantities. His quite reasonable conclusion is that QFT, as long as it relies on renormalisation to produce meaningful results, cannot be deemed a completely satisfactory framework for the treatment of electrodynamic interactions. The sentiment was also shared by those working in axiomatic approaches who perceived renormalisation as an indicator that QFT needed to be put on a more secure mathematical foundation. As we saw, several projects that sprang in this direction aspired to present a consistent language that would a) render the conceptual basis more mathematically accurate e.g. by carefully defining fields as distributions characterised by analogous axioms and b) define a theory consistent up to arbitrary scales.

As we noted, problems are not only of a mathematical nature. Physical insight into the very nature of the process as was described in chapter 0 is also lacking. A lot of the tools applied at a first level only possess a formal character and their physical interpretation depends on the broader framework adopted. For example, does QFT imply that spacetime is truly discrete (akin to a lattice) at some scale or should we opt for a more formal reading of the tools used – much like we do in dimensional regularisation? D. Fraser 2011 presents it in terms of a (representational) dilemma: assuming that the cut-off is real, i.e. corresponding to a breakdown scale associated with the discreteness of spacetime, then what sense does it make to take it to infinity after the completion of the procedure? On the other hand, if the theory works up to arbitrary scales (so that we may assume a spacetime continuum), then why do we need to introduce the regularisation procedure in the first place to obtain sensible results? Even if we adopt an instrumentalist reading of the cut-off, a slightly different question persists: why do so different regularisation schemes and renormalisation conditions lead to the same predictions? Is there an argument similar to what Schwartz 2014 (chapter 15) offers as an explanation of the independence of the Casimir effect from the regularisation method available for all cases?

All of these issues could simply be disregarded within a fully instrumentalist approach, but the very fact that they were taken seriously by several physicists at the time is a strong indicator of their deeper realist affinities. From a more realist perspective, something like the effective field theoretic approach originating with K. Wilson's work in the 70s was required to make sense of renormalisation in QFT. It is by analogy with condensed matter that we see why degrees of freedom at smaller distances may be neglected without harm. Since, at low energies high energy contributions will only appear in internal parts of the Feynman diagrams, they will become increasingly irrelevant for the infrared behaviour of the QFT system one studies. Thus physicists who acknowledged the importance of this puzzle correctly identified a fruitful question that could potentially (and, as it turns out, actually did) lead to a major conceptual breakthrough.

The important moral that we, as philosophers, can draw from this story concerns the role of a realist mindset in science. While not constituting a defence of the broader realist thesis, it does point to its fruitfulness for the interpretation and construction of scientific theories. Thinking along the lines of Stein 1989 on the contrast between instrumentalism and realism, we take the interesting part of the question of realism to concern the distinction between questions worth pursuing and those not-worth pursuing, i.e. to concern the role of theories as "resources for inquiry" (Stein 1989, p. 52). In this vein, identifying those parts of the theory that should be taken to play a representational role is far from being a straightforward business. Ultimately, the final verdict will be reached when one has deciphered their role in promoting advances in the process of scientific inquiry. For example, studies on the nature of ether in the 19th century (see Kragh 2011) as the material means for propagating light waves ended up being a rather fruitless line of research as were theories about the shape of electrons within a classical context. By contrast, questions about the definition of simultaneity, the stability of atoms in light of the Bremsstrahlung radiation, the existence of universality in statistical mechanics led to the development of techniques or even new theories, which, in turn, opened up whole new horizons in physics. Consequently, we view the contrast between instrumentalism and realism as a guide for theory construction as well as interpretation: which parts of a theory can be given an uncontroversial representational role, which ones are only to be seen as mere formal devices and for which ones are we best to suspend judgment in light of future developments¹.

¹This accords with Stein's central point that a "sophisticated" instrumentalism recognising the role of theories as tools for guiding the scientific enterprise need not be contrasted with realism. The thesis to be rejected is a rather naïvely defeatist anti-realism that reduces to stubborn skepticism about the nature of scientific results.

Note: There are cases in which the EFT program seems to fail. This is something to which we will return in the next two chapters, but a small comment is perhaps due here. Irrespective of whether the effective reading is right or wrong the merits of an at least partially realistic interpretive scheme will still hold. Even the demand for some physical insight (an explanation) into the success of so technical a procedure as renormalization is enough to ground the importance of an excursion to the questions surrounding realism in the narrower sense we gave it. And it is still worth extracting any useful ontological insights from the EFT program even within the, perhaps narrower, confines of success.

2.2 Effective Realism

We start with a brief discussion of the way realism is currently understood in the philosophy of science literature along with its central challenge, so-called the pessimistic meta-induction. This way, we wish to tighten the connection between the above historical remarks and the issue of realism, but also help motivate the transition to the more local, “effective” version of realism we will explore in the context of EFTs.

2.2.1 Motivation: Realism and Pessimism

The standard, or at least currently dominant, realist thesis is that scientific theories are to be read literally, i.e. as true descriptions of the world. Even in an approximate sense, theoretical statements featuring in explanations should be taken as true descriptions reflecting facts about the structure of the world while (certain) terms in these statements should be taken as referring to entities populating it. Thus, a statement like “the charge of the electron is x ” should be read “literally” as stating a fact analogous to “a page full of letters is in front of me” while the term “electron” refers to an entity existing “out there”, i.e. whose existence, much like that of the computer screen in front of me, is independent of my mental states (in contrast with the protagonist of the novel I was writing a few moments ago). The electron is endowed with a property, charge, which is also to be understood as indicating a true feature of the world and which can acquire an indirect meaning e.g. through the specification of a measurement procedure that will render the attribution of the property “charge” meaningful. Observe the apparent distinction between ordinary objects like tables and theoretical objects like the electron, which

cannot be observed macroscopically: the existence of the latter is to be decided indirectly usually through what is called “an inference to the best explanation” [see Psillos 1999, chapter 4]. The idea is that theoretical entities postulated in our best accounts of the phenomena are to be admitted in our ontology to the extent that they constitute indispensable referents of our theories: we need pions, charges, fermion fields etc to even state any predictions. The degree to which commitment to such entities differs from the way we are committed to objects of our ordinary experience (tables, flowers, spoons etc) is an open question (e.g. Van Fraassen 1980) and has sparked a lively debate with realists (Psillos 1999, chapter 9).

Of course, the problem is that the way science, and in particular physics, has evolved over the 20th century has shown that reading the ontology of the world off a theory is far from a straightforward business – contrary to what a “naïve” realist would presume. The main complication comes from the well-documented episodes of radical shifts in the course of scientific history in which newly emergent physical theories completely or substantially revised our picture of the world. An (in)famous argument in philosophy of science, the pessimistic meta-induction, turns this historical observation into an objection against the cogency of the realist thesis itself. At the core of this argument is the fact that highly successful physical theories of the past which shared key virtues (like fruitfulness, empirical adequacy, simplicity etc) with presently accepted theories, (which we now consider approximately true) were ultimately superseded by successor theories featuring a radically different ontology. What then warrants our belief that our current theories will not be revised in the same manner? If no reason to undermine the analogy can be presented, then all we can claim for our currently successful theories is that they are temporarily true. And truth with a provisional status can only be equated with falsehood – or not?

The typical realist response to this argument involves (a variation of) the so-called divide et impera strategy (Psillos 1999, chapter 5). In a nutshell, according to this defensive strategy, what we need to do to undermine the analogy is to scrutinise the incidents of theory change with an eye to distinguishing between those theoretical structures that were preserved in the succeeding theory and those that were abandoned. If we are able to show that the empirical success of the superseded theory was due to these shared, preserved rather than abandoned elements, we effectively create a defensive shield for the realist thesis: we will have found continuity within the apparent radical discontinuity. So, vindication of the realist thesis turns out to be a more “local” affair: essentially it hinges upon the success ratio of this strategy – the

final verdict to be reached in historical investigations of these transitional phases.

A vital component of this more localised version of realism is the acceptance of the fact that, strictly speaking, our theories never make true -fully accurate- statements about the world, but rather offer descriptions or statements that can only be taken as “approximately true”. It is clear how this can be used to safeguard against the pessimistic meta-induction: even if our theories might be rendered obsolete or revised in the future there is still “a sense” in which what they say about the world is correct. The local formulation allows us to check the realist thesis in a case by case manner by cautioning us to heed to the need to discern between parts of a theory enjoying empirical import and those lacking it. Note how this nicely aligns with the discussion of theories in the previous chapter. We noted that there is always an interplay between various layers of theory structure that dynamically evolve to meet specific needs in the description of phenomena. Accordingly, some parts of the theory are to be construed more formally as mathematical tools that will enhance the descriptive power of the theory. Identifying those parts admitting of a realist interpretation will often require a complicated story where physical insight will blend with formal tools to extract what is ontologically significant. As we are about to see, the renormalization group can provide hints towards this direction. The interpretative importance of these “hints” will prove to be far from a trivial matter, however.

2.2.2 Effective Realism: Definition

Since effective realism is a version of scientific realism, it is best we start with the standard characterisation of the latter. Typically, the scientific realist adheres to the following three basic commitments (Psillos 1999, Chakravartty 2017):

Scientific Realism [SR]

1. metaphysical: the world has a mind-independent structure, which scientific theories are meant to describe
2. epistemic: our best theories are able to provide us with knowledge about this structure both in its observable and unobservable aspects
3. semantic: scientific theories are to be read “literally” - their statements possess some truth value and their terms typically refer

The “effective realist” will not want to dispute the former commitment, which would imply some kind of idealism, phenomenalism, Kantian transcendentalism and so on. The world of QFT, of particle physics, even viewed through the lens of EFTs, will still possess a structure that will not depend on what epistemic agents like humans believe of it. A separate issue here is whether some levels of this structure become manifest without some form of intervention. This would imply some form of indeterminacy of nature at more fundamental levels - similar to a picture drawn by Adlam 2019². More on this possibility in the following (short) chapter. Even in this case, however, the kind of processes that will take place will not depend on anything more than the setup of the experiment and the same events will be assumed to occur under similar enough conditions, independently of the presence of humans. So, the metaphysical thesis must be considered secure.

The “effective” realist will, however, want to modify the epistemic and semantic theses to make them fit with their commitment to a more local version realism. Whether this modification can be done without sacrificing too much on the realist side is the challenging task ahead. When it comes to the epistemic thesis, the effective realist will need to say something about the way our best theories allow us to gain knowledge of the structure of the world and any limitations this procedure might involve. An important point emphasised by Williams 2019a is that, contrary to what standard accounts of theory interpretation usually require, in the case of EFTs we do not have access to a fundamental ontology. More specifically, we cannot tell what the world looks like according to QFT (at all levels of reality), simply because we can only treat QFT(s) as an accurate description(s) up to certain energy scales. Our epistemic access is thus restricted to some regimes, which can successfully be described even in full ignorance of small(er) distance physics. Consistent with discussions of patching in the previous chapter, theories are not closed systems from which we specify possible world models. On the contrary, the task at hand is to start from restricted theoretical frameworks consistently patched together and see how they might be applied to real world systems. The notion of approximation takes centre stage.

The semantic thesis will also need to be amended accordingly so as to reflect the special way in which terms refer to world structures in the EFT context. In contrast, say, with “classical” gas molecules, which can be construed or visualised as billiard balls bouncing with one another³,

²This could, in turn, bring us closer to some form of participatory realism akin to Fuchs 2017, but we will not pursue this point here.

³If ever such a “clean” picture was possible in classical mechanics to begin with! Part of the lesson of our

making sense of the way we employ the term “proton” is a more convoluted task. It is here that Tarskian semantics in terms of T sentences of the form:

*“the proton is attracting the electron” is true iff a certain proton is attracting a certain
electron*

ceases to be applicable in a straightforward manner. In a sense this is true of all theoretical statements that involve complicated truth conditions. The additional challenge in this context is that “protons” are picked out through an amalgam of facts about the various group symmetries applying, bound states of fields, relativistic constraints and of course triggered energy scales. Protons only emerge as stable formations at the lower energy scales at which quarks hadronise. At higher energy scales the term fails to refer to some intended particle-like formation. Part of the challenge for the effective realist is to offer appropriate conditions for endowing terms with a referential role that does not collapse to some form of empiricism.

All things considered, the effective realist thesis will be a modified version of the realist thesis of the following form:

Effective Realism [ER]

1. metaphysical: the world has a mind-independent structure, formed into semi-autonomous layers, which (individual) EFTs are meant to describe
2. epistemic: our best theories furnish us with approximate knowledge about this structure both in its observable and unobservable aspects within a scale of interest.
3. semantic: scientific theories are to be read “literally” - their statements possess some truth value and their terms typically refer to “stable” formations at a certain scale

ER is thus above all a moderate version of SR informed by the currently most successful meta-theoretical framework employed in physics. It is moderate in two ways. First, it is restricted to the framework of QFT. Consistent with variants of the *divide et impera* strategy to protect realism from the pessimistic meta-induction, it acquiesces to assessing realism on a case-by-case basis. It represents the maximal realistic commitments that can be salvaged within QFT as standardly practiced. For some, this might leave a bittersweet taste in mouth. All we

investigations is that we perhaps need to reconsider the unproblematic way we thought classical mechanics might apply - similar to the universalism we touched upon in the previous chapter.

seem to be doing is taking a stance, to examine QFT from a realistic standpoint, rather than argue for the latter's truth in virtue of its explanatory power in this framework. Consistent with our historical discussion above, however, we now see that one of the central merits of realism is that it can be of use in helping us sort deciding between fruitful or pursue-worthy from non-fruitle questions. All we want is some (local) epistemic criteria that will permit a distinction between the parts of the theory that we must deem contributing to the representational power of particular EFTs and those to be treated as purely formal tools.

The second way in which ER comes forth as a moderate thesis is, of course, through the caveats introduced in the standard version of SR. The important limitations in the kind of epistemic access and the semantics that we can employ within the EFT framework implies that typical realist claims about ontology or the way we refer to entities will need to be appropriately modified. Consequently, ER turns out to be not only a localised but also "conditionalised" species of SR.

2.2.3 Effective Realism: Prospects and Limitations

As Fraser. 2018 and Williams 2019a have both stressed, ER comes with the great promise of helping the realist basis key realist notions such as approximate truth and "selectivity" of ontology put on a more firm. This optimism has been put into question by Ruetsche 2018 (as well as Ruetsche 2020b) who not only doubts whether ER can be uniquely associated with realism –as opposed to some form of empiricism– but also raises skeptical worries on whether the EFT framework can play a guiding role for future physics. In this section we will begin examining these claims. We will see that while ER can indeed be of great help in refining the notion of approximate truth, it will be impotent to deal with the full implications of the problem of theory change. We will defer discussion of the empiricist leanings of ER to the next section.

(+) Approximate Truth Realists clearly need to avail themselves of some conception of approximate truth. It is vital for anyone wishing to claim that even though theories do not always offer entirely accurate descriptions of the phenomena or might be subsequently revised, they are still part of a line of progression towards a better understanding of the world. The scientific enterprise has an accumulating character with theories becoming more encompassing and more refined in spite of any (unavoidable) discontinuities that might arise in their development. For in what other sense Newton's theory of gravity, a remarkably successful theory widely used

today even in frontiers research in astronomy, can be said to be true in light of its being superseded by general relativity? By claiming that Newton's theory is approximately true, the realist can underline the fact that there was an important part of the theory that accurately latched onto the world and thus take GR to be "adding corrections" to it in certain regimes⁴. Presumably it will be possible to tell a similar story when a quantum theory of gravity supersedes GR. In this way one can still retain the key realist insight that one explains the success of theory based on its truth.

Now, there are two challenges for approximate truth. One: the notion seems to be too vague, metaphorical and unclear to be seriously considered as helping the realist cause. The realist seems to be trying to evade rather than confront the issue at hand. As Laudan 1981 has emphasised (p. 31-2):

Until someone provides a clearer analysis of approximate truth than is now available, it is not even clear whether truth-likeness would explain success, let alone whether, as Newton-Smith insists, "the concept of verisimilitude is required in order to give a satisfactory theoretical explanation of an aspect of the scientific enterprise.

Two: even assuming that approximate truth has been given a precise characterisation, it still remains an open question when (and how) it is applicable. To take an extreme example, how much of Aristotelian physics could be said to be an "approximately true" description in light of Newtonian theory? Or when is the NRQM description of stationary states adequate and when does it need to be corrected by QFT? We also know that the description of the same structures varies according to context. A coarse-grained description of a material or a fluid may not involve micro-constituents such as molecules, but rather treat as a continuous body. Similarly, the structure of a neutron or the details of weak interactions are suppressed in lower energy descriptions. Is it possible to reconcile the pictures?

Treating theories as effective descriptions within a certain regime can go a long way towards meeting these challenges while the RG can be used as a selective criterion for the ontological commitments one is supposed to make in QFT. When it comes to the issue of the multiplicity of descriptions⁵ EFTs offer a natural accommodating framework. As higher energy processes

⁴And indeed this seems to match scientific practice itself. Scientists and engineers will not try to produce accurate solutions from Einstein's equations, but will often rather try to correct a first approximation they obtained through Newton's theory.

⁵An issue one encounters in the context of classical theories as well. As Wilson's investigations in continuum mechanics have shown (M. Wilson 2017), in the description of materials such as steel one frequently comes across "descriptive conflicts" at the microscopic level which issue from the different dominant behaviours these materials manifest across the scales. We will have more to say about this soon.

are suppressed, heavy propagating fields will “collapse” or be reduced to local interactions. Further, the behaviour of systems at lower energies will exhibit some novel, semi-autonomous properties that a more fundamental theory cannot capture. This is the case with hadrons, for example: quark confinement, as Williams 2019a (p. 230) stresses, guarantees that at longer distances one will be unable to apply QFT without these formations. Thus, we know that at energy scales of the order of 10^2 MeV we expect to only see hadronic degrees of freedom. If quarks and leptons were discovered to be composite particles made out of preons, the autonomy of scales would imply that our current level of description would still be adequate for the kind of questions we have been using it.

Thus, the layered picture of reality emerging can accommodate seemingly incompatible ontological descriptions that are meant to be valid within some regimes. This helps vindicate the realist expectation that the replacement of one theory with one another can be realised in a continuous manner. The RG can offer clear guidance in this task through a quantitative assessment of the relevant contribution of higher-order effects. When the energy scales triggered are sufficient, for example, it will cease to be possible to neglect the propagating of smaller distance degrees of freedom such as the W boson. Another interesting aspect of the effective program is that one can hold on to competing theories as effective descriptions in anticipation of a successor completed theory. This is the case with gravity, which can be treated in a fully GR context or perturbatively as a low-order approximation to an underlying field theory.

(-) Two Challenges The success of ER in accounting for approximate truth cannot be extended as far as one might initially have hoped, however. Ruetsche in her recent work (Ruetsche 2018, Ruetsche 2020b) has flagged two sources of concern that at best severely restrict the optimistic outlook of EffSR and at worst threaten its very viability:

- **Effective Empiricism:** one source of worry is that the tools at the effective realist’s disposal can also be used by an empiricist. In brief, ER is committed to structures that remain robust under RG transformations like correlation functions or parametrising couplings. But the empiricist will also be content with any of these structures since they represent measurable quantities of the theory.
- **Pessimism:** another source of worry has to do with the capacity of RG techniques to address issues of theory change. In particular, it is entirely possible that a future theory like string theory will not be situated in the theory space spanned by the RG in QFT. Thus,

there is nothing that ER can tell us about the problem of pessimistic meta-induction.

We will start by discussing the second worry in this section and continue with the first worry in the following section. The reason is that while the second worry points to a significant restriction of the ER thesis, it does not present the kind of existential threat that a reduction of ER to some version of empiricism will imply.

Challenge 1: Theory Change While the RG and the effective reading of theories can render more precise the continuities between the various QFTs and their respective regimes of applicability, they have to remain mute when it comes to the connection between QFT and successor theories. It is true that treating the standard model as an EFT is a gateway to future physics. Taking the SM Lagrangian to be the first approximation in an infinite effective series:

$$\mathcal{L} = \mathcal{L}_{SM}^{(4)} + \mathcal{L}^{(5)} + \mathcal{L}^{(6)} + \dots \quad (2.1)$$

one can derive contributions from higher-order terms and check for deviations from the SM predictions. This will be a clear indicator that modifications to the SM are needed and perhaps new physics is around the corner. This is analogous to the way a theory, like the Fermi theory, can be known to be an insufficient description for phenomena above a certain scale: its breakdown indicates that a more fundamental description will be required. The way it works (see Petrov and Blechman 2015, chapter 8 for more details) is as follows. Impose the SM symmetry constraints on the higher-order terms of the above expanded Lagrangian. It turns out that for dimension 5 there is only one contributing operator, but dimension 6 operators in the so-called “Warsaw basis” is 64 (not including the extra complexity issuing from the different flavours). All of the terms can modify observables and low-energy constants (such as those governing top quark decays) that we measure experimentally and low-energy constants. If the predicted deviations are found, then this would constitute an indication that new physical degrees of freedom await above the electroweak scale of the SM.

Another way to probe beyond SM physics exploiting its effective status is to include expected or hypothesised contributing degrees of freedom explicitly in the effective Lagrangian. This includes weakly-interacting massive particles (WIMPs) that have been proposed as potential candidates for dark matter. If such particles interact with SM particles like the Higgs field, then their contribution can also lead to discrepancies between the measured values of certain observables within specific regimes.

Ruetsche 2018 has presented the following worry to the effective realist wishing to use the RG to make claims about theory change. Let's assume that Ω is the theory space spanned by the RG. Now, let \mathcal{T} be the low energy theory whose features we want to be preserved in a more fundamental, successor theory \mathcal{T}' . The central question is: does the space include all possible theories? On this basis, the skeptic can pose the following dilemma:

- Case 1: We have not fully specified the theory space Ω . Then, it is clear that we cannot say enough for \mathcal{T}' on the basis of \mathcal{T} alone. We simply do not possess enough information to eliminate possible alternatives – otherwise we would have done so.
- Case 2: We have indeed fully specified the theory space Ω . This might seem like a good starting point, but then what this space is missing is precisely the future theory \mathcal{T}' . And, it is perfectly reasonable to expect that \mathcal{T}' or some other \mathcal{T}'' lies somewhere beyond the region spanned by the RG.

In somewhat simpler terms, we can put forward the dilemma roughly as follows: On the one hand, many high-energy theories will be compatible with the low-energy structures we would like to preserve. Therefore, there is an underdetermination problem and thus RG techniques can only be of partial guidance to future physics. On the other hand, there is always the possibility that the future theory will be so “outlandish”, so radically different from the perspective of current meta-framework that the corresponding (RG) techniques will fail to apply. To make the latter possibility more concrete Ruetsche discusses the shift from Newtonian gravitation to general relativity: although the former succeeds in describing many phenomena to some degree of accuracy, assuming that some of the features involved in its explanatory accounts carry over to future theories would have led us astray. In GR we do away with concepts such central forces or action-at-a-distance. Even though “approximately true”, Newtonian gravitational theory fails to be a reliable basis for the successor theory.

But, let's focus on QFT for a moment. The problem arises when we think that the space of theories spanned by the RG in the QFT framework may have little to do with a future UV complete theory such as string theory. It is a remarkable fact about the string-theoretic description of interactions that they do not contain UV divergences⁶ (Witten 2015) and therefore one might claim that renormalisation is not required. It is in this sense that string theory is often

⁶The key difference is that interactions in string theory are no longer treated as point interactions, but rather involve strings extended in space. It is not free of the IR divergences we also encounter in QFT - but these can also be treated in the string framework (Sen 2015).

taken to be the(?) true fundamental theory of nature⁷. Given that string theory is a framework radically different than QFT, it is at best unclear to what an extent the RG can be a reliable guide for continuities between the two frameworks. Yes, as has frequently been noted (e.g. see the Weinberg-Gross exchange in Cao 1999), string theory should ultimately agree with QFT at low energy regimes –and therefore QFT will constitute an effective theory of ST–, but the effective realist will mostly have uncontroversial and trite advice to offer. This should come as no surprise: we know that multiple RG trajectories might converge to similarly behaving theories at low energies even though they will significantly vary in the UV. The EFT program specifically implies that the details about small distance processes will be hidden from macroscopic regimes.

So, where does this leave us? Most likely, a more restricted ER. The effective realist can use the RG to better track continuities between theories within the QFT framework: either by integrating out or matching one can put more precise boundaries on when certain effects will no longer be negligible. However, this is the most that one can achieve: future possible developments might lead to radical revisions⁸ for which the RG will simply fail to offer any hints whatsoever. As a result, ER is unable to cope with skeptical arguments of the historical pessimism espoused by the pessimistic meta-induction. This does not need to count as a threat to the ER thesis, however, since it was from the get-go construed as a local version of realism. Rather, the lesson is that we should rethink about the way a realist thesis should be construed in the QFT framework, as seen through the effective meta-framework.

2.3 Against (Some) Fundamentalism

The second challenge, which will be discussed in this section, will be the springboard for a treatment of a host of other issues having to do with the way theories, and EFTs in particular, model and give us epistemic access to the world.

Challenge 2: Empiricism The second challenge Ruetsche 2020b mounts against ER poses a direct threat to its viability as a realist thesis. The problem in a nutshell is the fol-

⁷Although the value of UV completeness as a guiding principle has been criticised in recent literature (Crowther and Linnemann 2019, 2017)

⁸And it's not only in the more ambitious project of string theory that radical departures might be taken. Even in less ambitious projects such as loop quantum gravity, substantive ontological shifts are involved: e.g. spacetime itself becoming an emergent structure.

lowing. If effective realist's commitments are to structures which remain invariant under RG transformations, such as correlation functions, then it be hard to demarcate ER from an effective empiricist thesis (EE), which restricts its commitments to what is accessible through observation (e.g. Van Fraassen 1980). More accurately, the empiricist, who is only interested in empirical adequacy than truth, will claim that ER specifically provides them with a means to eschew any talk about the metaphysical commitments of QFT and insist on what renders it capable of "saving the phenomena". In this case, of course, the phenomena will be something like the relative frequencies in scattering experiments.

There are two possible routes worth exploring here. First, we need to examine the kind of structures that ER has at its disposal for its ontological commitments. The task ahead in this case is to examine whether some of these structures implicate commitments that the empiricist will be hard-pressed to accord to. The second route concerns the (more?) intriguing possibility of disentangling fundamentalism and realism in an effort to preserve the latter even in the absence of an underlying fundamental ontology. We will note affinities between the case of EFTs and disciplines such as continuum mechanics or statistical mechanics (examined by M. Wilson 2017 and Batterman e.g. in Batterman 2001, Batterman 2013), in which we frequently rely on top-down models or regime-dependent descriptions to supplement or by-pass bottom-up modeling. By analogy, we will try to extend the kind of realist approach we take in those disciplines into the EFT domain. Apart from rendering fundamental ontology descriptions to some extent redundant, we will more interestingly see that, in accordance with continuum mechanics, effective descriptions possess an indispensable explanatory role within their respective regimes. Therefore, aiming at more fundamental descriptions will often be not only unnecessary but also undesirable or inhibitory.

2.3.1 First Route: Defining Commitments

Taking the RG seriously as a guide to ontology means that structures robust, i.e. invariant under RG transformations, are accorded special ontological weight and are expected to be preserved in future theories. The rationale behind this is quite appealing: since trajectories in the theory space converge close to the same point in the IR, then multiple theories in UV will exhibit very similar behaviour at low energies. Microscopic details will be largely irrelevant for the macroscopic behaviour we observe, but those parts of the theory that encapsulate the latter should be seriously taken as representational. What kind of structures are these, however?

Ruetsche discusses three possibilities:

1. Coupling Constants | 2. Correlation Functions | 3. Field-Particles

1. Coupling Constants, i.e. parameters such as the masses, interaction couplings or fields themselves are preserved under RG transformation in the sense that they represent physical quantities that will feature at low energies irrespective of the exact nature of small distance physics. Of course, their values will change depending on the scale under consideration, but once a particular scale is given, this value is fixed (on the critical surface) under changes at smaller distance physics.

2. Correlation functions are also a most natural choice since these are exactly the mathematical structures that the Callan-Symanzik equation shows to be invariant under changes of scale - with appropriate changes to the fields and coupling constants of the theory. Correlation functions are used to represent the propagation of fields and their interactions and be directly linked to scattering amplitudes through the LSZ formula. They thus capture the physical content of the theory. As an example, $\langle \phi(x)\phi(y) \rangle$ describes the propagation of a field from point y to point x while $\langle \phi(x)\phi(x)\phi(x)\phi(x) \rangle$ will indicate an interaction at point x . This seems to render correlation functions more favourable to a representational-realist leaning.

(Another candidate of the same kind could be the generating functional itself, which is also invariant under RG transformations. Apart from issues of well-definiteness (the generating functional should be as seen as describing a “procedure” rather than a well-defined mathematical object given well-known issues with the well-definiteness of the measure), however, a problem with the functional itself is that its physical significance is not as transparent as that of correlations functions - from which its representational role is after all derived.)

The main problem with the above suggestions, as noted by Ruetsche, is that they can easily play into the hands of the empiricist. The reason for this is that they are all quantities one can be measured experimentally: parameters like interaction constants are estimated by taking the ratios of different scattering experiments (like jet cross-sections for the strong coupling) while correlation functions are computed through their correspondence with scattering amplitudes. The empiricist will be interpreting things as follows. QFT is a powerful framework that allows us to obtain very accurate estimates about the frequencies of experimental outcomes we observe with an incredible degree of precision. However, this empirical adequacy is all we need to commit ourselves to: nothing compels us to make further claims about unobservables of

dubious status such as fields, virtual particles, radiative corrections etc. If anything, ER seems to induce the following conclusion: the stability of low energy structures against processes at high energy means that we can always restrict our commitment to a part of a theory that can be given an operationalist meaning⁹. The real underlying structure of the world remains as elusive as ever - contra any realist aspirations.

3. A natural response is to insist that these mathematical structures not only encode empirically accessible facts, but actually represent or, at the very least, map onto real physical processes and entities. What more natural entities to take than fields and their particle-like manifestations?¹⁰ There are a few issues with this realistic reading of the situation, however. First, unless some selective criteria are imposed, the aspiring realist will need to admit excessive ontological baggage. Ruetsche discusses the case of mirror fermions previously invoked by Williams 2019a as an example of the advantages of ER as a guide to ontology. Sure, she claims, the way mirror fermions appear in the Lagrangian changes with the way the effective description is realised but “the fact of their appearance does not”. Similarly, in the Lagrangians of non-abelian gauge theories Lorentz invariance is preserved through the addition of ghost fields. Should these also count as more than mathematical artefacts? If the negative answer relies on the fact that these fields do not come up in our measurements or low energy physics, we seem to have moved closer to an empiricist position - and likely abandoned our adherence to realism along the way. ER seems to be trapped between the Scylla of empiricism and the Charybdis of “over-commitment”.

Another issue emerging once we turn to an EFT reading of QFT is the status of the infinite terms showing up in the Lagrangian. In some cases, like in the expansion of ϕ^4 theory, these represent higher-order processes that get suppressed at lower energy scales. In others, like in the toy model of light scalar coupled to a heavy one, the Lagrangian might contain degrees of freedom corresponding to other fields that simply drop out at lower energies. Are we to extend the realist’s commitments to these fields as well? Ruetsche thinks that this poses a dilemma: if the answer is yes, then again we are being overly permissive; if the answer is no, then we seem

⁹No strict connotations should be associated with the term “operationalist” here. The main point is that we can always remain skeptical about structures that are not “empirically accessible”, that is, fail to be part of the mesoscopic world of our observation - measuring capabilities.

¹⁰There is a rich discussion about whether QFT supports a quasi particle ontology or whether it is fundamentally a theory about fields. Arguments both pro (e.g. Teller 1995) and contra (e.g. D. Fraser 2008) some particle interpretation have been presented along with no-go theorems such as those by Malament 1996 or Reeh and Schlieder 1961 – see discussion in Clifton and Halvorson 2001). For our purposes it will suffice to consider that, whether particle-like or field-like or x-like, there are indeed some features of the microscopic world that the abstract mathematical description of QFT corresponds to.

to agree on some graded notion of “reality”, i.e. fields will be more or less real depending on their contribution at a given scale. How can existence come “in degrees”?

The obvious move available, that is, granting the scale dependent nature of ontology, is construed by Ruetsche as a third puzzle: it conflicts with the way physicists typically think of other cases like phonons. The latter, collective excitations of more fundamental constituents in a crystal lattice, are (understandably?) not seen as on par with electrons dressed by virtual processes. Yet, if some significance is to be given to a scale-dependent conception of ontology, then phonons, as entities emerging at an effective level of description, should also be granted their graded level of existence. It is not clear that this should be of concern, however. As long as we can offer i) a realist underpinning to the terms we include in the Lagrangian and ii) a coherent sense of the “graded” nature of the reality, we will be able to see to what an extent the physicists’ view is vindicated.

2.3.2 Second Route: Reject Fundamentalism

The task ahead is to soften the above realisation, namely that underlying physical facts appear to be unparseable by a theory constructed along operationalist lines, i.e. almost exclusively tuned to measuring observable quantities such as correlation functions. To “massage” this difficulty away we will attempt to lend support to the following claims:

1. interpreting a theory in a realistic manner does not presuppose that a universal, complete or “straightforward” (a term to become clearer later in this section) description of the world be available. Such demands could end up being:
 - (a) Unnecessary since a lot of reliable information about the world can be obtained by theories that do not purport to offer reliable descriptions at all levels
 - (b) Pernicious since they might, as Williams (2019) puts it, disallow theories to “discharge their scientific duties”
2. our conception of realism should take on a more methodological tone in line with the main insights of our brief historical discussion in the beginning of this chapter.

Perhaps the major obstacle in appreciating the realistic vein latent in ER is the long-standing tradition of theory interpretation in philosophy. Theories should be analysable in something akin to a possible world semantics, with their models being coherent, maximally encompassing

pictures of the world. As we saw, the EFT program undermines this aspiration: QFT models can only fulfil a restricted representational role - one dependent on the intended accuracy and completeness of description desired. Now, this might seem hard to square with a realistic reading of theories since we do not seem to be going beyond a “recipe” for reproducing accurate experimental results. This is the attitude we would like to challenge here by separating realism from a commitment to fundamentalism. It is no small task to define the targeted view without committing a strawman fallacy. A tentative meaning we may attribute to this “tenet” is:

Fundamentalism #1: a scientific theory offers a set of basic constituents of reality, which (in principle) can be serve as the basis from which to (re)derive all physical phenomena in its intended domain

Thus, a Newtonian theory of mechanics would involve particles of matter interacting with forces and following trajectories in spacetime. Maxwell’s theory of electromagnetism would involve distributions of charged matter and force fields. In general, following a schema preferred by philosophers, a theory specifies n -tuples of the form: $\langle M, S_1, S_2, \dots, S_i, \phi_1, \phi_2, \dots, \phi_j \rangle$ with M usually representing the spacetime manifold or some abstract phase space and S_i -s and ϕ_j -s representing quantities either constrained kinematically or dynamically evolving. For example, in the case of GR one often writes possible models of the theory as: $\langle M, g, T \rangle$ with M the spacetime manifold, g standing for the metric field and T the stress-energy tensor. From the fundamentalist’s perspective, the ontological commitments of the theory

“...should be determined by assembling our various worldly claims into a grandly unified theory and surveying the amalgamated corpus for its sundry existential claims (= sentence of the form $\exists\alpha$ ”

M. Wilson 2017 dubs this the “Quinean thesis”, but I think it lucidly illustrates the way a “1-fundamentalist” take on ontological interpretation would look like. The mathematical apparatus of the theory is seen as being in some largely unproblematic (isomorphic?) correspondence with some basic elements of reality, which can be read off when the theory is presented in an appropriate form (e.g. axiomatised).

Now, it is arguably harder to find advocates of this view in the philosophy of physics literature, but not so hard in the area of (analytic) metaphysics where a lot of claims are made about fundamental constituents of the world in terms of dispositions (e.g. Mumford 2002) or the mosaic view à la Lewis (e.g. Loewer 1996, Lewis 1994). Thus, one might find people like

Sider arguing for the possibility of the metaphysics project in “physical vacuo”, i.e. above and beyond the final verdict of first-order discussions of physics, along the following lines:

And just like scientists, metaphysicians go on to construct general theories based upon those observations, even though the observations do not logically settle which theory is correct.[...] Observation bears on metaphysics in a very indirect way, and it is far less clear how to employ standards of theory choice (like simplicity) in metaphysics than it is in science. [...] You don’t need to have answers to all meta-questions before you can ask first-order questions.

Needless to add, effective approaches in physics speak against metaphysical projects for which the fundamental nature of reality can be accessed simply by flying over the trenches of our furthest physical advance into the territories of the unknown. But, of course, this is hardly an interesting target for philosophers of physics who have been more circumspect in their tackling of ontological questions. The idea that science provides (progressive) approximations to reality has been in the air since at least the time of Newton.

More interesting “fundamentalist” targets can be identified in standard discussions of quantum ontology. Ruetsche 2020b, for example, mentions Maudlin and Wallace as two of an “abounding” number of fundamentalists. However, it is not clear how we can count proponents of so diverging projects under the same interpretational banner. Clearly, some disambiguation is called for. We will try to decipher Ruetsche’s claim by defining the additional notions of fundamentalism - one that can be associated with interpretational lines closer to Maudlin’s and one closer to Wallace’s. Our view will be that the effective meta-framework undermines the motivation for adopting a primitive ontology approaches as well as “overly trusting” readings of a theory. Let’s take on each one in turn. ¹¹.

¹¹Here is some background to render the whole analysis more self-contained: Discussions in the foundations of quantum mechanics have frequently centered around the reality and nature of the quantum state. Within the realist camp, i.e. those taking the quantum state to represent something physical (in contrast with more instrumentalist or epistemic views that accord the state only some calculational or agent-specific role), a debate has sprung over the following question: does the wavefunction suffice as the basis for complete system descriptions or does the theory need to be supplemented with some additional structures or, what has been called, some “primitive ontology”. Typically, the Everettian camp (e.g. Deutsch 1997, Saunders 1993, Wallace 2012, Carroll 2019) are sympathetic to quantum monism: the quantum state of the multiverse is all we need to recover all observed aspects of reality. By contrast, the Bohmian camp (Dürr and Teufel 2009) postulates particles with (unknown) determinate positions whose trajectories are determined by their pilot-waves (wavefunctions). These fundamental mass distributions constitute the primitive ontology of the theory. Proponents of collapse theories are divided, with some arguing in favour of an additional primitive ontology in terms of flashes or matter densities while others, like Myrvold 2019b claiming that this is both redundant and problematic – a reading one could also read into Ghirardi, Grassi, and Benatti 1995.

Fundamentalism #2: a scientific theory requires a specification of a clear ontology [an answer to the question: what is “out there” in space and time] and a dynamics that describes how this evolves over time

Note that this view does not claim that the primitive ontology will be fundamental in the #1-fundamentalist sense. This means that no one denies that the wavefunction (or the EM field, Newton’s particles etc) might be further reduced to some more fundamental physical structure as revealed by successor theories. The fundamentalist claim here is different: it is requirement for the well-definiteness and “explicitness” of a theory’s ontology. Local beables, the “out there” structures constitute an indispensable precondition for any sensible physical theory. Here is Maudlin 2007 pondering about the way physics would be (or perhaps could not be) like without local beables:

First, it is rather hard to see why a theory that lacks local beables altogether would bother to postulate anything like [1] spacetime: after all, if there is nothing in any local region of spacetime, why think there is a spacetime? Furthermore, if the local beables and the locations are removed from the physical ontology, it is hard to see how [2] evidential contact with the world is to be made except at the level of conscious experience. Local beables also make transparent the explanation of the [3] intersubjective character of physics.[...] What is there outside of the various observers that all the observers could independently become aware of, and hence agree on?

Certainly, the idea that a theory essentially consists of some ontology and some dynamics for its evolution is neither outlandish nor unmotivated. Au contraire! From the very conception of a mechanistic philosophy we find the idea that everything in the universe can be described in terms of some extended substance that moves (evolves) in accordance with some laws. As Allori 2013 remarks (p. 62) “As in classical mechanics, it seems most convenient to explain, if possible, the behavior of familiar macroscopic bodies postulating that they are composed of microscopic entities in three-dimensional space that constitute the fundamental building blocks of everything else”. This is an approach that appears to have worked well in 18th mechanics and even 19th century electromagnetic theories: postulate some ontology like fields or particles and track their evolution through the appropriate dynamical laws. It is hard to resist the temptation of elevating this scheme to a precondition for a well-defined theory. And this is exactly what Maudlin 2007 appears to think when he writes:

But one might also try instead to derive a physical structure with the form of local beables from a basic ontology that does not postulate them. This would allow the theory to make contact with evidence still at the level of local beables, but would also insist that, at a fundamental level, the local structure is not itself primitive. [...] This approach turns critically on what such a derivation of something isomorphic to local structure would look like, where the derived structure deserves to be regarded as physically salient (rather than merely mathematically definable). Until we know how to identify physically serious derivative structure, it is not clear how to implement this strategy.

The EFT construal of theories renders obsolete the idea that ontology is something that must be declared “from the outset” as a precondition for the well-definiteness of the theory. Instead, we saw that, much like in Curiel’s example for the black hole at the center of Milky Way, the ontological commitments of a theory, the story it tells us about the world is something to be recovered through the intricate process of matching theoretical claims and reality. By following Disalle 1995, Wallace 2012 (ch. 2 & 3), Myrvold 2019a), we find it sufficient if a theory can be used to accord with our observations, produce novel predictions and reconstruct those aspects of reality that we seem as falling within its domain of applicability. This means that the link between an object like a ball rolling on a plane and its microscopic description in terms of wavefunctions could remain fuzzy without loss of coherence or understanding.

Note that the criticism laid out here is of an external rather an internal kind. The claim is not that there is an inconsistency or internal coherence in a primitive ontology (or similar) project. It also does not imply that when completed, those approaches abiding by this “dictum” be they Bohmian mechanics, collapse theories or even many-worlds theories equipped with some primitive ontology will not furnish a deeper or more intuitive understanding of the quantum. In fact, it might turn out that arguments such as those presented by Allori 2013 and Maudlin will be vindicated through practice: projects taking seriously this (currently) heuristic demand will be led to significant conceptual breakthroughs¹². Nevertheless, from an external point, the motivation to reject or modify the currently most widely adopted framework in order to adhere to some aprioristic maxim mostly hinges upon how compelling it is from a pragmatic, i.e. motivation-related perspective. In light of the effective program such a modification is not forced on us.

It would equally be a mistake to ignore the intricacies of theory-reality correspondence by

¹²For example, one could take work by Okon and Sudarsky 2016 to be sympathetic to such arguments.

becoming too trusting of a given formalism. This can be seen as another form of “fundamentalism”, one that hastily seeks to dissolve conceptual issues or impatiently desires an interpretation of a theory at “face-value”. This goes against a more mediated approach to interpretation which accords with the previous chapter’s discussion:

Fundamentalism #3: taking a theory “seriously” requires a strictly literal interpretation, which amounts to a (direct) mapping from the formalism to “elements of reality”

Proponents of views such as wavefunction realism (like Albert 1996, Ney 2019) or Everettian interpretations of QM argue that the most natural, “straightforward”, conservative or attuned-with-practice understanding of QM is to take the quantum state as representing something existing in (some) space and evolving with time. For the former camp this implies that we must accept that what quantum mechanics tells us is that the wavefunction is to be treated as a field “living” in a $3N$ -dimensional, where N corresponds to the number of particles in the whole universe. Just as the electric field is distributed throughout $3D$ space and assigns values to each of its points, so the wave-function is multi-dimensional field assigns values to this superspace of $3N$ dimensions. A most straightforward reading of the formalism! The Everettian camp is not committed to this admittedly bizarre ontological picture¹³, but is also (in most variants) a monistic view sharing a deeper methodological affinity with wavefunction realism. The key difference (at least in the more recent versions of the view) is the role assigned to decoherence: the histories that decohere, or for all practical purposes neatly separate from one another, are treated as the basis for the emergence of a huge multiverse of quasi-classical domains resembling separate worlds (slightly different versions of reality). This interpretive conservatism of Everettian interpretations is usually presented as a virtue. For instance, Wallace 2012 writes (p. 38):

The ‘Everett interpretation of quantum mechanics’ is just quantum mechanics itself, ‘interpreted’ the same way we have always interpreted scientific theories in the past: as modelling the world.

and a few pages later (Wallace 2012, p. 43) continues:

To be sure, alternative positions of a sort are available. One can lapse into instrumentalism. One can try to find some way of making sense of the theory that is not exactly instrumentalism, but that nonetheless does not take the theory at face value. And one can try to find

¹³Wallace 2017 in particular has criticised it in various occasions as ill-defined and unmotivated.

an alternative theory that is as explanatorily and predictively successful as quantum mechanics but whose interpretation is more to one's liking [...] However, one can do this with any scientific theory that one dislikes. [...] The fact remains that, at present, there are no known ways of explaining the quantitative predictions of relativistic quantum theory other than Everett's. The Everett interpretation is the only game in town.

However, as Halvorson 2019 has remarked the word "realism" seems to be doing a lot of heavy-lifting in this context. Not only does it prescribe that we take the theory seriously, but leads to the most extravagant ontological picture, the strongest arguments for which are a) the idea of a "literal reading" of a theory and b) a form of TINA argument against any rival program (Wallace 2020, p. 98):

Does it set out to reform the practice of quantum physics, and does it provide evidence that this is more than a bluff by actually doing the hard work in some non-trivial, concrete examples across multiple instantiations of the quantum framework [...] If the answer in each case is 'no' . . . , well, maybe don't hastily commit it to the flames, as it may contain valuable insights and be the seed of a yet-to-be- completed research programme, but don't kid yourself that it is at present a viable interpretation of quantum mechanics, and maybe be a bit cautious exploring all its metaphysical implications until you've done some more work to see if it plausibly might be made viable.

But none of the two arguments really has bite, especially in light of all we have seen in this chapter. First of all, the project of taking a theory at face value crucially depends on what formulation of the theory one takes as the basis for interpretation. And, there are at least two complications with this. On the one hand, a lot of what is treated as the standard formulation is the outcome of historical contingencies. For example, if the de Broglie-Bohm approach was consistently formulated in the 1930s when the foundations debate was raging on one might be dealing with a different sort of theory as the textbook presentation¹⁴. Even if one responds that we should focus on QM as standardly practised by physicists, it far from obvious what this frame would be. Quantum mechanics, as Wallace himself claims, is more like a framework that differently adopted by condensed matter physicists, cosmologists, information theorists, chemists depending on their purposes. That the unified ontological underpinning of all these aspects is an emergent multiverse is far from trivial.

¹⁴Needless to say, different questions would have been prioritised in this case too: perhaps a lot more mental power would have been focused on relativistic extensions of QM.

On the other hand, it is not the case that a theory can be interpreted literally (at least as this term is understood by Wallace) in an unambiguous manner. True, there might not exist any compelling reason to treat quantum mechanics differently than all our previous physical theories (although this is certainly ambiguous), but even in successful theories of the past it is not obvious what formulation to use for the ontological implications (e.g. Lagrangian vs Hamiltonian formulations), what a theory is telling us about specific domains¹⁵.

Before closing this section, we briefly turn to what we could call the pernicious restrictiveness of the fundamentalist mindset. What is pernicious about it is that it often compromises the explanatory depth of a theory (Williams 2019a p. 229). Insisting on an “in-principle” bottom-up derivability is wrong-headed: the computational intractability of the problem would probably foil any serious attempt for a clear solution and emergent phenomena would be ignored. Now, it is possible to claim that it is only practical limitations that render this problem unsolvable and “in principle” one should expect the behaviour of hadrons to be reduced to that of quarks. However, one would be hard-pressed to put more flesh on the bare bones of this counter. How does this enhance our understanding of the world and how does it help us with expanding our theoretical framework? If anything it appears to hinder progress in the latter (by leading to pointless digressions) and fails to appreciate the top-down nature of some constraints for the former (as the following two cases will reveal).

2.3.3 Classical Analogues

To alleviate any suspicion of contrivance, it will be instructive to consider some exemplary cases for which fundamentalism of any sort ceases to be relevant. Perhaps remarkably, we do not have to go too far; we can restrict our search to the classical setting and examine a) the way material scientists describe metals at different descriptive levels using modeling units that are mutually incompatible with respect to the microscopic structure they postulate and b) top-down approaches to problems in statistical mechanics.

Case 1 M. Wilson 2017 is a rich source of philosophical investigations into the intricate ways in which scientists formulate and adjust theoretical frameworks so as to ensure that they are applicable to the intended target system and computationally tractable. Among Wilson’s main lessons is that in their bid to achieve these goals, scientists are frequently forced to adopt

¹⁵See the relevant “debate” on the foundations of classical mechanics between North 2007 and Curiel 2014 as to the status of the different formulations such as the Lagrangian vs Hamiltonian.

descriptive schemes that make conflicting claims about a system at different levels of its organisation. Drawing on examples from material science, Wilson systematically exposes a problem known in the literature as “the greediness of scales”, which he characterises as follows:

What is this difficulty? Each RVE scale-focused modeling will utilize differential equations, and their descriptive demands inherently reach down to the infinitesimal level. But these demands often differ. The result: under amalgamation direct descriptive conflicts will arise in the same vocabulary with respect to the properties that steel displays on small-scale levels.

The problem identified here is that the description of the structure of a metal in terms of differential equations at a certain scale will require an assumption A that will conflict with an assumption B required to describe it at a different scale. However, both representational models also require that their assumption holds to arbitrary small scales – to ensure that the mathematics will be well-behaved. We are thus unable to obtain a coherent picture of the microscopic details of the system. The result, as Batterman 2013 puts it (my emphasis):

...even though we often have good models for material behaviors at small and large scales, it is often hard to relate these scale-based models to each other. Macroscale models represent the integrated effects of very subtle factors that are practically invisible at the smallest, atomic, scales. *For this reason it has been notoriously difficult to model realistic materials with a simple bottom-up-from-the-atoms strategy.*

Another interesting strategy frequently employed, which is intriguingly analogous to what one comes across in EFT construction, is to input empirical or macroscopic information (in the form of constraints, boundaries or omission of terms) in an attempt to make the equations involved computable. For example, if one knows that a metal will bend after pressure it is applied, they will incorporate this information by adjusting the parameter encoding fracture strength. “In this manner, multiscale techniques represent a deft compromise between the purist top-down and bottom-up methodologies that once divided nineteenth-century philosophy of science into warring philosophical camps” (Ibid, p. 224). In perfect agreement with the gospel we have been preaching so far in this chapter we read later on (p. 227):

The outputs of multiscale modelings supply mixed-level explanations in the sense that their descriptive architectures generally stem from direct empirical observation of the manner in which various RVE scales causally affect one another within a complicated material.

We don't pretend to have "derived" these empirical hierarchies from molecular fundamentals; we instead exploit our direct knowledge of physical layering to better computational advantage.

Note the similarity with the situation we find ourselves in QFT: even in the allegedly unproblematic context of classical physics, sometimes all we can hope to achieve is a more contextualised, scale-dependent representation of systems in complete ignorance of the underlying microscopic structure. Should this be taken as an argument against the representational capacities of these models? I think that this would be an unnecessarily radical conclusion to draw: all this forces us to do is reappraise the scope and function of scientific realism to harmonise it with practised techniques of scientific modeling.

Case 2 Similar lessons can be drawn in the context of statistical mechanics. Batterman has long insisted on the importance of top-down modeling opportunities in the study of phase transitions (e.g. Batterman 2001, Batterman 2013, Batterman 2017). With the development of renormalization group techniques, it has been possible to describe the behaviour of thermodynamic systems around critical points (where there is a phase change like the transition to a paramagnetic or ferromagnetic phase in a magnet or vapor phases in fluids). The fact one exploits in this case is that system fluctuations around these points are so dominant that they effectively "wash out" any intermolecular forces, which would furnish information about the microstructure of the system. Thus, it becomes possible for systems of very different underlying constitutions to exhibit the same macro behaviour - a fact known as universality. While parameters required for an accurate description of macroscopic phenomena are in principle connected with the underlying microscopic structure of the particular system, deriving their values is virtually impossible in a completely bottom-up fashion (Batterman 2017). Dissenting philosophical treatments of these phenomena, such as that by Butterfield 2011, have emphasised the pragmatic, i.e. more instrumentalist, role of these techniques. The key point, reminiscent of Ruetsche's worry in the QFT context, is that lacking some microscopic description the proper stance to take is that of instrumentalism: while these techniques form powerful tools for computational purposes (the good behaviour of mathematical structures), they do not necessarily carry physical significance or at the very least not the kind of physical significance accorded to them by people like Batterman.

Irrespective of whether one wants to go all the way along with Batterman's conclusions

about the role of emergence in phase transitions, I think that the above considerations help motivate the following, more modest, position. Consistent with our discussion of scientific theories in the previous chapters, we acknowledge that the vastly intricate way in which modeling is achieved implies an equally intricate and involved task of distinguishing between representational and formal parts of a theory. What the cases studied show is that top-down modeling or the blending of top-down and bottom-up modeling techniques can often be the only viable means of producing descriptions for systems whose excessively complex constitution rules out pure bottom-up modeling. Within their domains, these models are typically treated as representing the behaviour of a system at a given scale even if their assumptions are known to fail when extended to lower regimes (e.g. hydrodynamics). Note that this does not alter any commitments we make with respect to the microscopic. We still believe sea waves to be composed of water molecules, atoms and quarks even if we cannot derive their behaviour from them – or if we have to treat the systems as continuous to ensure the applicability of specific mathematical tools.

Morals Let us wrap up this discussion by returning to Ruetsche’s concern. The dilemma put forward was that ER’s tools are either susceptible to empiricist appropriation or ontologically too permissive. Now, the above cases help motivate a sense in which we retain a realist-leaning understanding of the effective program. First, the realist needs to abandon a commitment to fundamentalism as a precondition for exploring the ontological implications of a theory. An EFT description confined to an appropriate regime of applicability can still be treated as representationally significant in a manner analogous to the other cases of physics examined. The RG can be a vital tool in characterising the contributions of relevant processes within each specified regime. Even when a more fundamental theory is missing, the matching procedure can be used to input “by hand” the suitable terms for a given scale of description. Second, the cases above also clearly show that EFTs are not singular in this respect. The blending of top-down with bottom-up modeling observed in the study of materials, continuum mechanics and statistical mechanics, attests to its ubiquity as a representational strategy. It is not only in EFTs that one finds a fine interplay between adjusted parameters and experiment to obtain models aligned with experience.

When it comes to the issue of “graded reality”, I think that the more pragmatic dimension hinted at by Ruetsche 2020b towards the end of her paper is promising, but, appearances

notwithstanding, need not be construed as going against a realist position. Recall that what the worry is: even if suppressed, the processes are still “there” even as terms in an infinite Lagrangian. It seems paradoxical to say that their reality comes in degrees. One can potentially counter this by adopting a structuralist metaphysics: recovering the ontology from dynamical patterns e.g. in the way argued for by Wallace 2012 (chapter 2 & 3) might be turn this into a more cogent proposition.

However, even if ones does not want to pursue this way out, it is perfectly possible to be non-committal about certain ontological and representational aspects of the theory without precluding the possibility of an interesting and useful realist thesis. In fact the Ruetsche’s moderate empiricist, who disavows variants such as constructive empiricism, can hardly be said to be an interesting rival for a more moderate realist position:

The humble empiricist differs in significant ways from more notorious empiricists. To my mind, these differences all redound to the favor of the humble empiricist.

Here are some of these differences:

[1] Some non-realists (e.g., Cartwright 1999) contend that the physical world is irremediably untidy or irreconcilably disjointed—not the sort of thing to afford the kind of truth conditions that would vindicate fundamentalism. The humble empiricist, by contrast, entertains the possibility of a true, fundamental theory T_{final} . [...]

[2] Empiricists are often characterized as believing that the success of science requires no explanation. This wrongs the humble empiricist, who offers an explanation for T’s success. It’s not the explanation the fundamentalist favors. [...]

[3] Such underdetermination makes agnosticism about T’s account of hidden springs permissible: T succeeds not because it is the unknown truth T_{final} but because it mimics the truth at scale \updownarrow . [...]

To my mind the first two points already blurry significantly the boundary between the empiricist and the realist. Perhaps we find ourselves in a similar spot to that envisioned by Stein 1989 in his treatment of the scientific realism debate during the 1980s. A sophisticated instrumentalist, here corresponding to Ruetsche’s moderate empiricist, may have no real bone to pick with a more flexible realist. A more pragmatic identification of structures bearing representational weight can equally well serve a more methodological brand of realism. The core of this realist variety will be the need to separate those parts of a theory that play a formal

(assisting) role (like the alteration of the spatiotemporal dimension in regularisation) and those providing some physical insight about the world (e.g. the anti-commutation relations obeyed by fermions¹⁶).

Does the proton exist? Do quarks exist? Yes, in a manner that is not dissimilar to the way objects of ordinary experience exist. They serve a particular role in a conceptual framework: they codify structures featuring in causal processes that ultimately reach our measuring apparatuses. Perhaps the way we employ and refer to them in the QFT context is more convoluted than, say, in classical mechanics, but is this a reason to be skeptical or non-committal about their existence? Perhaps we do not need to abandon the realist pretension to uncovering the structure of the world, but only a narrow-minded, fundamentalist conception that puts special constraints on the representational function of theories. It is in this respect that treating phonons as disanalogous to quasi-particles (e.g. electrons dressed with a “cloud” of virtual excitations) seems ill-motivated. The basis for such a preferential treatment needs to be firmer than that “in the case of phonons, according to our best picture of the underlying physics, there are no particles there; in the case of dressed electrons, there are particles there, just not quite the particles featuring in our effective physics” (Ruetsche 2020b, p. 313).

Finally, when it comes to the issue of permissiveness, it is true that more work needs to be done. The problem is that the cases examined above are meant to be in the context of non-fundamental theories, something that differentiates them from QFT. Ruetsche’s criticism of Williams’ mirror fermions example shows that perhaps the RG is not a sufficient tool to pick out the structures meriting ontological significance. Still, the fact that at we can have a common low-energy basis to form expectations about the space of structures that any underlying high-energy theories should agree on is an important first step in helping us navigate through the unknown lands of BSM physics. Irrespective of that, it is no less important to be more transparent about the kind of realist thesis we have in mind. This is what we will turn to next.

2.4 How Realism Matters

Here is the deal. The currently dominant variant of realism makes such strong requirements that a realistic understanding of QFT, and quantum theories as they currently stand, is rendered problematic, if not impossible, from the outset. Even when any fundamentalist inclinations are

¹⁶In the sense that they were not just a mathematical trick to introduce a lower energy bound, but had implications for the very behaviour of these quanta - the kind of statistics they obey.

excised, the connection between quantum reality and our knowledge of it remains loose. The operationalist character of textbook-like presentations of QM is inimical to a realist view aimed at providing the referential basis of theoretical terms. Wave-functions, quantum states, operators, operator-valued distributions etc are employed for the sole purpose of obtaining accurate predictions of experimental results, while their representational content is left unspecified. Accordingly, it is no uncontroversial task to present an interpretation of what an operator, like the spin operator represents, whether quantum states should be seen as encoding information about some physical structure “out there” or about an agent’s knowledge and so on. The challenge is further aggravated by the effective reading of QFT which makes the connection between the empirical and underlying ontology even less direct. Therefore, if a realist reading of QFT as both a quantum and effective theory is to be possible, we had better come up with a more relaxed formulation of realism than the received view.

A realist approach to a scientific theory should at the very least be characterised by the following features:

- **Quasi-representationalism** The theory (loosely) tells us something about the world.
- **Minimal Coherence** Its predictions cohere with those of other established theories within the same domains of applicability.
- **Fruitfulness** The concepts introduced lead to novel results and/or create expectations about unexplored phenomena.

So, even if we cannot treat a quantum field in the straightforward way we can treat a classical field because a quantum field is, to follow Teller 1995, a “determinable” as opposed to a “determinate”, we should still be in a position to acknowledge that QFT has significantly enhanced our knowledge of the micro-world and revealed novel facts about it. It would be a defeatist stance of utmost degree to accord all the claims about the types of fundamental fields and forces, the connection between spin and statistics, the symmetries constraining possible interactions, the separation of scales and more a mere fictional status.

To illustrate this point, let us consider QED. Even though the measurement problem is as unresolved as ever and the meaning of all its formal tools has not been fully explained, there are still important facts that QED is (or appears to be) revealing about the world. For instance, the interaction term $e\bar{\psi}A^\mu\psi$ implies that photons do not couple with one another [quasi-representational]. The number of possible polarisations of A^μ [two] also agrees with

facts known in the classical setting [minimal coherence]. Similarly, facts such as vacuum polarisation or the running of the coupling constant led to a QFT novelty, namely the notion of “dressed” particles, which down the line led to the development of the renormalisation group [fruitfulness].

I believe that the three above desiderata comprise the nucleus of a realist thesis. Rival anti-realist views will dispute at least some of these claims. Constructive empiricists will probably wish to deny any representational-referential status to theoretical terms used. They will also disregard the potential capacity of concepts to go beyond their reach and produce novel predictions – as all one is interested in is empirical adequacy. The instrumentalist of the logical empiricist kin will probably remain neutral with respect to the referential status of theories, might take coherence to be something like a regulative principle (the unity of science!), but will also disregard the surplus content of theoretical terms. The (minimal) realist, however, will definitely need to accord to some referential or representational capacities to theoretical terms and agree that they can be relied upon to guide one’s expectations about presently unexplored phenomena or even future theoretical shifts. Coherence with other theories or frames is also an indispensable demand; if there is a (determinate) mind-independent structure of the world then it would be peculiar, bordering to absurdity, to have certain regimes described by mutually incompatible frameworks¹⁷. Let us summarise some broad perspectives on realism in the following table:

¹⁷The project of making sense of the connection between statistical mechanics and thermodynamics can be thought of along these lines. If the underlying microdynamics fails to reproduce the macroscopic behaviour we have failed to create a coherent picture of reality; the puzzle pieces do not match.

Table 2.1: A Realism Typology

Anti-realism	science is one practice among many we may adopt in our engagement with the world
Instrumentalism	a theory should be taken as a tool for systematising and reproducing empirical results
Methodological	treating aspects of a theory realistically can lead to important conceptual breakthroughs
Truth-based	the standard realist thesis one normally associates with some theory of reference – multiple variants of this view exist

The methodological kind of realism suggested here has an affinity to what Saatsi 2020 has recently coined “progress realism”, which, as a local and minimal form of realism emphasises the representational capacities of theories without being committed to a specific answer with respect to the reference of their terms:

According to this realist tenet, theories of mature science, such as quantum theories, latch onto unobservable reality in ways that are responsible for their empirical successes—both predictive and explanatory — as recognized by scientists. Progress realism is perhaps a good label for attempts to defend this realist tenet.

and later on continues:

We can note, as a purely conceptual point first of all, that in defending this tenet a progress realist is not making an assertion about the world, or about what we can claim to know. [...] Indeed, it is possible that one is only able to argue that an appropriate representational relationship holds, without being able to tell exactly what that relationship is like.

The methodological kind of realism espoused here is in accord with this more relaxed take on the representational relationship between theory and the world – as is evident from the rejection of a more “literal reading” of theories à la Wallace. However, it is specifically aimed at tackling and illuminating this connection using a multiplicity of (perhaps “indirect”) tools such as:

- **no-go theorems** force a choice among incompatible assumptions; this has proven important in the debate over hidden-variables or the nature of the quantum state in QM

- **conceptual unloading** eliminate assumptions and principles or challenge their status (play the sceptic) until something falls off the tree; the following two chapters can be said to follow this methodological vein
- **historical tracking** reveals the role played by certain theoretical posits or mathematical structures in the development of the theory; the work of Stein 1967 on Newtonian space-time could be seen as an instance of this
- **rational reconstruction** (often combined with historical undertakings) seeks to optimise the theoretical framework identifying the right or minimal assumptions necessary¹⁸ to characterise it
- **practicality** often necessitates revisions to the formalism in order to address computational limitations in the treatment of complex systems; Wilson's case studies can be seen as attempts to unearth the ontological lessons of non-fundamentalist theories

These are only a small sample of the techniques one can use to comprehend the reasons for which some theories are better representational devices than others: no representation without deliberation. Or, for archaeophiles: there is no royal road to ontology!

There are several advantages to adopting this version of realism. First of all, it sets aside more or less redundant debates between empiricists (of the kind described by Ruetsche) and realists over the global stance one should adopt with respect to EFTs. In particular, worries of the constructive empiricist kin are more properly met at the more abstract, philosophical level where one can challenge the presumed voluntarism of the empiricist (irrespective of the particular content of a theory). Therefore, the question of whether correlation functions, for example, are open to empirical appropriation becomes moot. The methodological form of realism developed here seeks to re-orient the discussion to questions that truly matter for the betterment and furtherment of our scientific theories. Unfortunately, as is the case with most exploratory endeavors, there is no a priori guidance as to what these questions should be. Frequently, all we can do is realise post factum that a certain question led to breakthroughs (compare Dirac's fascination with large numbers to dissatisfaction with earlier renormalisation procedures).

With this more local and situated form of realism we free ourselves from the need –as well as the vain temptation– of reading theories in ways that correspond to our philosophical

¹⁸To be clear, it is hard to always disentangle the two “tactics”. In fact, for the purposes of extracting ontological lessons it would probably best to reconstruct historical episodes so as to be rid of superfluous structure. This has been a central goal of philosophical treatments of space-time theories.

predispositions. Realism turns into a relevant issue, an important methodological device for our treatment of each physical theory. In fact, we may speak of an inference to the currently best methodology (ICBM): is it fruitful to split our theory into a part to be taken seriously as representing something “out there” and a part that is to be seen only as an instrument to facilitate derivations? Much like the way Stein 1992 construes the Carnap versus Quine debate, the issue at hand is not whether one view can be proven incoherent, which, although always an efficient way to cut down on the number of alternatives, is a rather tedious exercise in conceptual bookkeeping. The shortcut route (and arguably the most persuasive) is to evaluate its effectiveness in describing and prescribing effective scientific practice. As long as the quest for realist underpinnings of a theory is a fruitful one, the strategy of adopting a realist stance is vindicated.

Naturally, the question now is how can draw this realist-instrumentalist line within a specific theory. The short answer would involve theory construction: the representational capacities of the theory, the way it latches onto the world and the way its terms-structures acquire referential content requires paying attention to the specifics of all the strategising we talked of in the previous chapter. We need to appreciate how a given theory is adapted to a particular set of phenomena, how it is frequently cut around the edges to adjust to contexts of failure or how it is modified in order to cooperate with some independent theoretical framework. This is something we can appreciate in the way Stein 1989 attempts to dissolve the realism vs instrumentalism debate and the way he analyses the misgivings of Poincaré’s methodology in (Stein 2021):

The basic mistake that I ascribe to Poincaré is that of seeing the significance of theoretical work as residing essentially and exclusively in its function in organizing knowledge (putative as well as real): that is, organizing the “real generalizations” – which count as presently claimed knowledge.

At the heart of his critique lies the idea that the goal of theoretical terms is not only to systematise facts about our experience (as, to engage in some well-intended caricature, an instrumentalist of the positivist kin would have it), but rather to actively stimulate the expansive character of the scientific enterprise. This means that, in agreement with Newton’s methodological dictum¹⁹, we should try to extend the scope of our conceptual scheme, barring any inconsistencies,

¹⁹Especially rules of reasoning III (“Those qualities of bodies that cannot be intended and remitted and that belong to all bodies on which experiments can be made, should be taken as qualities of all bodies universally) and IV (“In experimental philosophy, propositions gathered from phenomena by induction should be considered either exactly or very nearly true not withstanding any contrary hypotheses, until yet other phenomena make such propositions either more exact or liable to exception”) (Newton 2014).

as far as empirical results allow. Indeed, comparing Poincaré to Einstein, Stein writes:

And this is the crucial difference, as I see it, between Poincaré's relation to the special theory of relativity and Einstein's. Both of them discovered this theory—and did so independently. So far as its mathematical structure is concerned, Poincaré's grasp of the theory was in some important respects superior to Einstein's. But Einstein “took the theory seriously” in the sense that he looked to it for NEW INFORMATION about the physical world – that is, in Poincaré's language, he regarded it as “fertile”: as a source of new “real generalizations” – of empirically testable consequences.

Between the two, only Einstein put enough (realist) faith in the theory to create expectations about subsequent investigations. It is, to use an parable of sorts, like a group of lost adventurers, who desperately looking around on a deserted island, are fortune enough²⁰ to uncover a map in some abandoned temple. Only those brave souls that decide to bet on the accuracy of the map's depictions are truly committed to its representational potential. Those unwilling to try their luck along the path that map indicates as the one leading to salvation clearly do not (anti-realists) while those that find its successful record a mere coincidence bear close similarities to the (constructive) empiricist. The methodological realist would try out the map, risk within a low-stake context and then raise the stakes as successes accumulate.

Remark: A final word on the realism typology before wrapping up this section. Obviously, the classification here is not and does not claim to be exhaustive. One can further fine-grain the taxonomy by including variants like perspectival realism, pragmatist realism, structural realism, constructive empiricism, methodological anarchism and so on. Each of these approaches has its own idiosyncrasies that would merit a more detailed examination than can be offered here. However, for our purposes, suffice to note that I take the four identified types to be the archetypal stances one can take with respect to realism – with the above variants being refinements or sub-species of the four main views identified. For example, structural realism would be close to the truth-based realism trend while perspectival and pragmatist realism would be placed on either side of the methodological version. An interesting question is to explore whether these three formulations essentially bog down to more or less the same thesis – especially within the quantum frame.

²⁰Note how “luck” nicely corresponds to what realists have called “epistemic luck” or what has been swept under the rug here, namely, how the theory manages to fulfil its representational role.

2.4.1 EFTs: An Application... or Three

Perhaps it is worth seeing how the brand of realism sketched above plays out in the EFT context. In particular, we want to see how some of its central tenets such as the representational capacity of certain structures, the instrumental value accorded to others and the open-endedness of concepts (which is reflected in the guiding role of unresolved questions) can be said to better track the rationality of some specific cases within the EFT framework.

Example 1: Fermi Theory Fermi theory provides (again) a most convenient case with which to illustrate some of the above points. Recall that, according to our modern standpoint, the theory is a low-energy description of the electroweak theory essentially “collapsing” the propagation of bosonic degrees of freedom (such as in the τ decay) to a 4-point interaction (graph taken from Burgess 2020):




Figure 2.1: Propagator to Point Interaction

Fermi theory is a very clear instance of the way in which “guiding questions” led to a substantial revision and refinement of the extant theoretical framework. For it was noted early on that the theory was not predictive at energies above 100 GeV and thus non-renormalisable (for its infinities could not be cured to all orders). The search for a renormalisable theory was eventually vindicated by the development of the electroweak theory. The EFT framework accommodated this insight by showing quantitatively what went wrong with the operator behaviour in the Fermi theory. As long as energy scales were low enough the propagation of the W boson:

$$\frac{ip^\mu}{k^2 - m^2 - i\epsilon} \tag{2.2}$$

simply collapsed to a point interaction when $p \approx M_w^2$. In this case the propagating momentum of the boson is simply negligible compared to its mass and the whole interaction can be

described as governed by a coupling constant:

$$\frac{G}{\sqrt{2}} = \frac{g^2}{8M_W^2} \quad (2.3)$$

Consistent with the above take on realism, the EFT framework implied that the breakdown of the theory: i) flags the regime of applicability of the Fermi theory and ii) signals the existence of a more fundamental theory underlying these processes. On the other hand, the descriptive content of the theory is to be taken seriously, i.e. as representing processes within the energy regimes in which the theory produces meaningful arithmetical results. The existence of a more fundamental theory does not contradict or belie the representational role of the Fermi theory. In fact, the EFT framework allows us to perform “smooth” transitions from one theory to the other in that it lets us anticipate when certain effects – processes can be neglected and when they become dominating – indispensable.

Example 2: Chiral Perturbation Theory Lest the above example gave the impression that the question over the ontological commitments of a less fundamental theory can be resolved within a more fundamental theory, let us consider the case of chiral perturbation theory. As is well-known, perturbative methods in QCD break down at low energies. As a result, it is impossible to directly go from a description in terms of quarks and gluons to a description in terms of the low-energy degrees of freedom, namely, mesons and baryons. Essentially, one needs to match between theories that contain different degrees of freedom. Fortunately, one can use the EFT techniques to circumvent this problem and extract a low-energy description for QCD by relying on the following phenomenological tactic.

This idea is to construct an effective Lagrangian which will be consistent with the symmetries of the underlying theory, i.e. QCD. At the massless limit, the QCD Lagrangian for light quarks u and d written as the following doublet:

$$q(x) = \begin{pmatrix} u(x) \\ d(x) \end{pmatrix} \quad (2.4)$$

turns out to be:

$$\mathcal{L}_{QCD} = -\frac{1}{4}G_{\mu\nu}^\alpha G_{\alpha}^{\mu\nu} + q_L(i \not{D})q_L + q_R(i \not{D})q_R \quad (2.5)$$

which is symmetric under $SU(2)$ rotations for q_L and q_R separately. This chiral symmetry²¹, $SU(2)_L \times SU(2)_R$ is broken when masses for the quarks are included:

$$\mathcal{L}_m = -m_q(\bar{q}_L q_R + \bar{q}_R q_L) \quad (2.6)$$

The key to moving beyond this purely QCD description is symmetry breaking. When the temperature of the universe dropped enough to meet the Λ_{QCD} scale it went through a phase a symmetry breaking phase transition:

$$SU(2) \times SU(2) \rightarrow SU(2)_{isospin} \quad (2.7)$$

with the low-energy degrees of freedom consisting of isospin invariant pions. The “trick” now is to introduce matrices combining the pionic degrees of freedom π^+ , π^- , π^0 into a matrix of the form:

$$U(x) = \exp\left(\frac{i}{F_\pi} \begin{bmatrix} \pi^0(x) & \sqrt{2}\pi^-(x) \\ \sqrt{2}\pi^+(x) & -\pi^0(x) \end{bmatrix}\right) = \exp\left(\frac{i}{F_\pi} \sigma^\alpha \pi^\alpha(x)\right) \quad (2.8)$$

with σ^α the Pauli matrices. One then goes to construct the most general Lagrangian consistent with the unbroken Chiral symmetry:

$$\mathcal{L}_\chi = \frac{F_\pi^2}{4} \text{tr}[(D_\mu U)(D_\mu U)^\dagger] + L_1 \text{tr}[(D_\mu U)(D_\mu U)^\dagger]^2 + \dots \quad (2.9)$$

The constant F_π^2 is “controlling” the strength of processes of pion interactions such as $\pi\pi \rightarrow 4\pi$. One can also include source terms, which help describe decays of kaons and pions. One can further generalise to an $SU(3) \times SU(3)$ by replacing $\pi(x)$ with $\phi(x)$ such that $U(x)$ end up being:

$$U(x) = \exp\left[2i \frac{\pi^\alpha T^\alpha}{F_\pi}\right] \quad (2.10)$$

so that the Chiral Lagrangian can be used to describe bound states of three quarks, i.e. baryons. Remarkably, using spontaneous symmetry breaking, it is possible to be ignorant of the exact mechanism that produces these low-energy condensates (the excitations we detect as baryons), but still construct effective theories that are descriptive.

²¹Actually the Lagrangian is invariant under a $SU(2)_L \times SU(2)_R \times U(1)_V \times U(1)_A$ symmetry with the latter two associated with vector and axial currents.

We note that despite the fact that we possess an underlying fundamental theory, EFT techniques allow us to renounce the strict requirement of deriving the low-energy, or less fundamental, ontology from that of the deeper theory. Yet, this poses no threat to the “reality” of either theoretical level nor to any realist aspirations we might have: pions, protons, neutrons are as much “real” or representationally significant as the quarks that compose them. Treating one theory as an effective description of the other allows us to maintain “bonds of cohesion” between the two regimes (knowing when one theoretical description breaks down or is intractable and need to switch to the other) even without explicitly reducing, in the stricter sense of deriving, one to the other. Far from being a limitation, the effective description in this case allows us to extract what is representationally significant in an otherwise impenetrable regime. In fact, historically, the model also led to the prediction of the Ω^- particle – significantly enhancing the belief that we had gotten the story right.

Example 3: Beyond the Standard Model Research in BSM physics is particularly well-adapted to this methodological species of realism. Indeed, physicists determined to expand high energies physics beyond the narrow confines of the SM have been treating the SM Lagrangian as an effective theory, i.e. akin to a first order approximation of a more fundamental theory which comprises higher order processes:

$$\mathcal{L}_{fn} = \mathcal{L}_{SM} + \mathcal{L}^{(5)} + \mathcal{L}^{(6)} + \dots \quad (2.11)$$

$$= \mathcal{L}_{SM} + \frac{1}{\Lambda} \sum_k C_k^{(5)} \mathcal{O}_k^{(5)} + \frac{1}{\Lambda^2} \sum_k C_k^{(6)} \mathcal{O}_k^{(6)} + \mathcal{O}\left(\frac{1}{\Lambda^3}\right) \quad (2.12)$$

where $\mathcal{L}^{(n)}$ includes terms suppressed by the appropriate energy scales Λ^n . Attempts at expanding the SM proceed as follows: include new (admissible) terms such as:

$$Q^{(5)} = \epsilon_{jk} \epsilon_{mn} H^j H^m (l_p^k)^T C l_r^n \quad (2.13)$$

which is consistent with gauge symmetries and called “Weinberg’s operator” (with “C” being the charge conjugation operator). What is remarkable about this term is that it predicts the violation of lepton conservation – a feature not shared by any other term in the \mathcal{L}_{SM} . Unsurprisingly, then, this creates a clear expectation for what processes we should be able to detect as we probe higher energies in particle colliders. Terms such as the above lead to corrections in the probabilities of processes, so that any experimental discrepancies from current results will

be indicators that BSM physics has been found.

Adopting the more flexible branch of realism suggested here legitimises our “taking seriously” what a theory like QFT is telling us, even in its rather patchy and semi-operationalist formulation. We are thus in position to acknowledge that the theory is capable of guiding our expectations about what processes are realisable in light of accepted constraints (in terms of symmetries, conserved quantities or desirable features such renormalisability) and how these might be revised in light of new evidence. At the same time, it is also capable of guiding us ahead by suggesting possible revisions (elimination of some symmetries, inclusion of extra of terms etc) that can be tested as possible predictions of a modified theory. In this sense, a “reflective equilibrium” can be seen at work here: from principles to expected results and vice versa. Methodological realism is proposed as the relaxed form of realism that can accommodate this uniqueness of QFT(s) especially in light of its(their) effective status.

2.5 Conclusions

Let’s sum up this rather long discussion. Realistic aspirations are a good motivation for identifying fruitful problems and re-constructing, revising and expanding our theories. Thus, formulating a sensible realist thesis in the context of EFTs is not an irrelevant task. Effective realism seeks to transform and improve on the realist position in accordance with the main lessons gained by the effective program. Examining modern attempts at clarifying and defending this thesis, we saw that they met with considerable success when they define it in a local manner, i.e. adapted to the QFT-as-EFT framework – as opposed to a general thesis about all science. Still, we realised that even this is not enough. An empiricist sceptical counterattack forced us to abandon some ground and retreat to a more secure, but perfectly satisfactory realist thesis that we branded as “methodological”. This approach blends the standard realist and instrumentalist views with an emphasis on uncovering those aspects of theories that should be viewed as representationally anchored to the world as opposed to mere formal (e.g. computational) tools. The RG can serve as a tool for a preliminary assessment of which structures are expected to be taken seriously and which ones not, but it is not enough. To be able to separate the wheat from the chaff, one will usually need a combination of methodological and historical considerations: how the theory connects with previously established theoretical frameworks, i.e. how well-understood terms and concepts map or reduce to or change into new ones, how its domain of

applicability was delineated and how experiments or observations guided and constrained our stance against particular structures (think of Shimony's experimental metaphysics). Of course, none of this is in any way particular to EFTs – a fact we take as a further indication that the effective approach to theories can serve to refine, rather than revolutionise, our methodological tools. We will turn to a more focused treatment of the ontological implications of EFTs in the following (interlude) chapter.

Before closing this chapter, however, it is worth adding a few remarks on the relationship between idealisations and realism. *Prima facie* the use of idealisation seems to run contrary to realist sympathies. When we engage in idealisation we present an (intentionally) “false” picture of the target system with the intention of simplifying the task. This, however, immediately gives rise to the following concern: how can general principles such as dynamical laws, symmetry constraints etc be taken to say something “true” about the world? Cartwright 1983, but also Van Fraassen 1980 have pressed this point. The laws of nature are in the final analysis “false”. For Cartwright, in particular, the real explanatory weight is on the various causal capacities of the systems – with abstract laws offering some scheme-like patterns applying only to very idealised situations.

It is not hard to see that our preceding analysis of effective realism as well as of the structure theories in the previous chapter renders such a die-hard instrumentalist view otiose. For, in the quasi-hierarchical structure we presented for theories, laws and highly idealised descriptions of systems have a pretty distinct constitutive role at the upper levels of the stratified theoretical framework. We saw that symmetry principles (such Lorentz invariance), for example, can be constitutive of a whole framework and essentially correspond to the adoption of a specific perspective on the way we mean to describe physical systems. We saw that choosing the framework and then the appropriate dynamical description does not, as is to be expected of course, directly lead to an accurate, correct or even adequate representation of the target system. The situated kind of reasoning implicit in more pragmatic approaches will be needed to extract the empirical content of the theory. In this sense, the laws of nature are not some kind of heavenly dicta that we come to confirm or dis-confirm by directly checking with experience. Rather, they constitute something like (road) signs of inquiry and are therefore highly entangled with actual scientific practice. It is this through the continuing success of their frameworks that they come to be accepted or revised. Rejecting them as false *propositions* is to treat them in a manner that, in light of actual practice, they were not meant to.

Chapter 2b

Interlude: The World according to EFTs

What kind of ontological picture emerges once we take the EFT program seriously? A paper by Cao and Schweber 1993 has sparked some controversy about the lessons one should draw for the ontology, epistemology and methodology of physics. Cao & Schweber take the EFT program to have shown that fundamentalism, the quest for a final theory, is a dead end and thus we should consider seriously the idea that all we have is an infinite tower of EFT up to all scales. The paper has drawn considerable attention as well as critical replies. We will analyse their view and evaluate its cogency in light of these criticisms. However, an important enterprise for us will be to explore the possible ontological interpretations one could give to EFTs. Essentially this depends on the stance one takes with respect to EFTs. There are roughly four attitudes in the literature:

Table 2b.1: Attitudes to EFTs

Enthusiasm	Cao & Schweber
Optimism	Crowther, Fraser J, Williams, Wallace
Caution	Butterfield & Bouatta, Ruetsche, Hancox-Li (?)
Scepticism	Fraser D, Halvorson & Mueger

The Sceptics We have already seen that D. Fraser has extensively criticised the practised version of QFT for its apparent lack of mathematical rigour and its failure to coherently combine SR with QM. We also saw that she views renormalisation and the renormalisation group techniques as more or less formal tools through which one extracts the empirical content of the theory, but which do not have to say anything about its theoretical content. Thus, it comes as no surprise that, according to her, we may not base our ontological commitments on the EFT

program. Following this line of thought further one might be tempted to say that this is only a temporary framework, one we employ out of convenience, up until a consistent fully rigorous version of QFT emerges in agreement with tenets of some axiomatisation or “rigourisation” program. A closely related sceptical sentiment towards standard QFT is the natural endpoint of Halvorson & Mueger who also insist on the need to present within some mathematically rigorous framework. Clearly, the “skeptical” road is the path of the patient: in the absence of a completely satisfactory theory, thou shalt not make any ontological commitments.

There are two main “complaints” that can be levelled against this stance: i) it fails to consider any lessons that follow from the currently best theory available for subatomic phenomena on rather “idiosyncratic” philosophical presuppositions and ii) more generally, adopts a stringent methodological criterion that fails to properly engage with scientific practice. Leaving the second point aside for now¹, it is somewhat disappointing for philosophers to refuse to engage in conceptual issues that emerge in physical theories (e.g. issues about inconsistencies and the role of mathematical tools such as asymptotic expansions, guiding principles such renormalisability or naturalness and emergence in light of EFTs) just because they fail to conform to some standards about what constitutes an interpretable framework from a philosophical perspective.

The Optimists The recent interest in EFTs and practiced QFT has moved philosophers away from the more stringent demands of the received views about theory interpretation to more flexible approaches, more favourable to the EFT framework. The optimist camp in general sees the conceptual advances of the 1970s as not only a path towards a better understanding of renormalisation, but also a rethinking of the way we approach QFT and physical theories in general. The scale dependence of physics and the breakdown of particular QFTs at certain energy levels reveals something deeper about the very structure of the world and signals the uniqueness of QFT as a theory. That the theory, or better, that each theory in this framework, more or less announces its own regime of applicability is an indication that a more fundamental description is presumably missing. This is justified on the basis of a theory’s behaviour under RG transformations: as one moves to lower energies only certain processes will be relevant, with those corresponding to more complicated interaction diagrams essentially being suppressed. What is thus the picture of the world that follows from this optimistic perspective? It is quite appealing in fact! Theories may only be consistently applied within certain regimes to produce accurate

¹We will return to this in the last chapter where we will discuss the more methodological implications of EFTs and especially their role in guiding research in future physics.

predictions. At their breakdown scales, some successor theory will be needed. However, as long as we restrict ourselves to degrees of freedom appropriate to the regime we investigate² we can rest assured that we will obtain a consistent and adequate description. Since different layers decouple from one another, the complicated processes of fundamental levels will remain inconsequential for any macroscopic descriptions of interest. QFT thus presents us with layers of semi-autonomous regimes, which, much like the coarse-grained descriptions of statistical mechanics, will hide their detailed structure as we move towards larger distances.

The optimist can remain agnostic as to the possibility of a fundamental theory, which will replace QFT, or the “tower of EFTs” scenario. Usually, certain inadequacies of QFT are considered clear indications that some theory will need to replace it: failure to convincingly incorporate gravity, consistently describe black holes (as well as early universe cosmology) and consistently reach high energies (given the eventual breakdown of QED at its Landau pole). One can take the key lesson of EFTs to be that the standard model is also an EFT and will thus require some modification and/or eventual replacement by a more encompassing theory. Interestingly, this need not be taken as invalidating the ontological lessons we have managed to draw from currently used QFT. Having explicitly acknowledged its “effective” status, we anticipate that present day descriptions will be approximations to those of a deeper theory³. The optimist of this kin will likely be quite well-disposed to the reductionist program permeating the spirit of the particle physics community (as described by Weinberg 1987): penetrating further into the more fundamental constituents of the world we are bound to enhance our understanding and the explanatory power of our theories.

The optimist is confronted with two main challenges. First, how to think of this layered picture of reality in connection with a possible UV complete theory. Can we consistently hold on to this picture of progressively revealed structures when we find out that there is a minimal length scale like that associated with string theory? Another way to put this is: can the optimist escape the fate of turning into an enthusiast⁴? Apart from this, the optimist will also want to address a more pressing concern: how to think of apparent breakdowns of EFTs in cases

²The relevant degrees of freedom are not arbitrarily chosen, of course. An effective expansion will give an indication of the processes that will be relevant for a particular scale. In a bottom-up approach, choosing what terms to include will frequently depend on a physicist’s insight and experience.

³To clarify a bit, as this might seem in tension with some of the claims in the previous chapter: the approximation in mind here is of a more informal kind. It highlights the typical requirement that a more fundamental theory will need to reproduce the results of QFT at sufficiently lower energies.

⁴As we shall see, there are reasons to resist this extra step. However, the question here from an ontological point of view is whether the more reserved optimist approach can be coherently maintained.

such as the Higgs boson mass or the cosmological constant? Do these problems imply some restrictions for the EFT program, do they point to some new interesting physics ahead or are they themselves products of confusion?

The Cautious A more inhibited stance is to accept that the EFT program has indeed contributed to a furthering of our physical understanding of QFT, but try to be more circumspect about its scope. According to this, more cautionary, approach, one accepts that that EFTs a) have or at least can shed light on the process of renormalisation, whose success on Butterfield & Bouatta's (2014) words seemed like "manna from heaven" and b) can put inter-theoretic relations on a firmer ground. Reservations in embracing the optimist's viewpoint can come from the possibility of an instrumentalist or empiricist reading of this program. We have already seen Ruetsche 2020b present a potential empiricist appropriation of the effective realist's commitments. The cautious might also seriously consider the apparent problems with the Higgs boson mass or the cosmological constant as indicating inherent limitations of EFTs; thus rejecting its universality, especially with respect to phenomena involving gravity (e.g. Koberinski and Smeenk forthcoming).

Evidently, the cautious must also be patient. Since they are reluctant to join the optimists in fully embracing EFTs as a universal meta-framework to think about theories, but also wish to refrain from the sceptics' insistence on a revised framework for QFT, they also need to (partially) suspend judgment about its broader ontological significance. While they might avail themselves to RG tools in accounting for the success of renormalisation, they will resist treating QFT (fully) as an EFT (e.g. Li 2015) or question the kind of morals we can draw from such a move (e.g. Ruetsche 2018, Ruetsche 2020b). On the other hand, the cautious can be an optimist with a provisional status: they might accept that the ontological picture drawn by the optimist is correct only for a restricted range of theories within the QFT framework. Perhaps other areas such as gravity fail to adhere to its basic tenets and require a wholly different approach (again Koberinski and Smeenk forthcoming). On the other hand, problems such the hierarchy problem might lead to a further refinement of the notion of autonomy in the context of EFTs (more on this in the following chapter).

The Enthusiasts Among the views positively predisposed towards the EFT program, the most radical approach is of course the one endorsed by Cao and Schweber 1993 (p. 71-2):

The EFT approach extends the atomistic paradigm further, and within that framework the domain under investigation is given a more discernible and a more sharply defined hierarchical structure. [...] On the other hand, taking the decoupling theorem and EFT seriously would entail considering the reductionist (and a fortiori the constructivist) program an illusion, and would lead to its rejection and to a point of view that accepts emergence, hence to a pluralist view of possible theoretical ontologies. [...] More precisely, what is to be rejected is the suggestion that it is possible simply by means of these kinds of connections to infer the complexity and the novelty that emerge at the lower energy scales from the simplicity at higher energy scales, without any empirical input.

Truly subscribing to the EFT program, for them, entails that one also commits to an abandonment of the idea of a fundamental ontology. Just as probing higher and higher energies leads to the disclosure of further and further structure (e.g. with more and more higher-order processes occurring or virtual electrons polarising the vacuum), our theories will also reveal a progressively richer underlying structure as we try to describe physics at ever smaller distances. Instead of arriving at *the* fundamental level at which spacetime itself would be discretised or strings would vibrate, we would simply replace one fine-grained description with another. Probing is a process going on ad infinitum. The result (Ibid, p. 66):

We thus obtain an endless tower of theories, in which each theory is a particular response to a particular experimental situation and none can ultimately be regarded as the fundamental theory.

This anti-fundamentalist view not only seems to contradict our most cherished methodological practices of the past, when reduction reigned supreme, but also appears to make little sense in light of the ontological picture of the world we have been accustomed to since the mechanistic turn of the 17th century. How can there be no end to the infinite layers of processes across energy scales? How should we think of all these “veiled” occurrences? If anything, it is disturbingly reminiscent of a “it’s turtles all the way down” attitude! One possible picture is to adopt something like what we might call a “top-down enforcement” that we find in Adlam 2019. Using the concept of probability and objective chances for motivation, Adlam raises the possibility that the term “fundamental” might signify “an admission of defeat”: “We question as deeply as we can, but eventually we grow tired, plant our flag in the ground...” and might in fact be based on a false (reductionist) expectation that “things would get simpler as we got further down, and eventually we would be left with an ontology so simple that it would seem

reasonable to regard this ontology as truly fundamental and to demand no further explanation”. The following, in particular, seems remarkably close to capturing the kind of ontological picture enthusiasts of the EFT program wish to draw:

But one might argue that this is getting things the wrong way round: the laws of nature don't start with little pieces and build the universe from the bottom up, rather they apply simple macroscopic constraints to the universe as a whole and work out what needs to happen on a more fine-grained level in order to satisfy these constraints. Presumably at least some features will be left underdetermined by the global constraints, and that is where the arbitrariness comes in, but there is nothing wrong with this as long as the arbitrary features are of the harmless kind.

Now, in many respects the enthusiast seems to be saying something analogous. Since the tower of EFTs is inexhaustible, there is no end to the fine-graining procedure and since no fundamental theory exists there is no way in which one can “build the universe from the bottom up”. The QFT framework provides a set of broad constraints (especially symmetry principles) that should apply to all theories formulated in it. Consistent with the ER thesis, low-energy matters of fact will be fixed irrespective of small distance processes. Since physics at these levels will be largely irrelevant for obtaining accurate low energy results, one might be tempted to claim that that part of reality will be hidden, perhaps even undetermined in the “harmless way” envisioned by Adlam. This convergence to similar low energy physics might help motivate the more radical conclusion that the world only determines these high energy structures once probed to do so. In this case we might be faced with a rather intriguing prospect:

...as we build bigger and bigger particle accelerators to probe ever more deeply, the universe will be forced to invent deeper and deeper levels of reality that exist only to answer our questions.

To what an extent is this ontological picture accurate, however? It is never easy to dampen someone's enthusiasm, but there seem to be compelling reasons to do so in this case. First, as Hartmann 2001 has stressed, the conceptual basis of Cao & Schweber's argument is shaky. For one, to make their case of a tower of quasi-autonomous domains, Cao & Schweber rely on the decoupling theorem of Appelquist and Carazzone 1975. According to this, when we examine two coupled fields, one of which is heavy (high energy) and the other light (low energy), it is possible to include the effects of the former to the latter through a readjustment of parameters.

This is the essence of decoupling and something we have seen at play when integrating out heavy fields to obtain an effective Lagrangian. A crucial assumption of the theorem, however, is that the theory we start with, and on which we perform these operations, is renormalisable. Of course, it is possible for Cao & Schweber to simply jettison this argumentative line and insist on the bottom-up construction of EFTs similar to the “extreme version” discussed by Castellani 2002 (p. 263):

The EFT approach in its extreme version provides a level structure (“tower”) of EFTs, each theory connected with the preceding one (going “up” in the tower) by means of the [renormalization group] equations and the matching conditions at the boundary...

The problem in this case is that there is no guarantee that this program will indeed continue to work in the manner needed by the enthusiast. It merely seems to beg the question against the proponent of a final theory.

Apart from problems we touched upon in the cautious stance, there are indications that an infinite tower of EFTs cannot be the whole story. There have always been discussions about the possibility of a minimal length or some form of spacetime discretisation from the early days of QFT (see Hagar 2014, chapter 4-5), but various candidates for a quantum theory of gravity seem to suggest radical departures from our standard conception of spacetime. Even if a framework like string theory is wrongheaded, alternatives such as loop quantum gravity or causal set theory heavily dispute the continuity of spacetime. It is unclear how the ontological picture delineated above will square with a departure from the typical continuous four-dimensional manifold picture.

The above concerns notwithstanding, a more methodological point should also be considered. Oftentimes in the history of science the full implications of a theory become known once a successor theory has been found. This can arguably be said to be the case with our understanding of spacetime in a Newtonian context after the development of the theories of relativity. The process of uncovering the unique features of the latter also led to a clearer and deeper appreciation of the kind of spatiotemporal structure implicit in the former. Similar remarks can be made in optics or the development of electromagnetism and ether. Thus, in the context of QFT too, it will be imprudent to fully embrace the apparent implications of the current framework so as to exclude the very possibility of a UV complete theory. In this sense, we can endorse the more pragmatic take on the EFT program expressed by Crowther 2016 (p. 79):

We should recognise that EFT is not really a “program” that requires die-hard subscription

and fights in opposition to the search for a final theory, but rather a practical necessity, and wonderful aid to progress in physics. The success of EFT does not mean there is no final theory... and viewing EFT as effective means remaining open to all possibilities...

We can summarise the ontological picture adopted by each attitude towards EFTs in the following table:

Table 2b.2: Ontological Temperaments

Enthusiast	Tower of EFTs; no final theory; fundamentalism is wrong; the further higher energy scales we trigger the more structures we reveal – ad infinitum
Optimist	Autonomy of scales; QFT valid up to certain energy levels; each level might be dominated by different structures; a final theory will potentially emerge to describe the most fundamental level
Cautious	Potential insight; the RG can teach us something about the success of renormalisation and the connection between processes and energy scales; unclear how broadly can we trust its lessons
Sceptic	Mathematical tool; progress in understanding the world according to QFT will only come from a completion of some axiomatic program; RG etc are non-rigorous techniques

2b.1 The (Preliminary) Verdict

We have reasons to take the effective reading of QFT seriously. First, there is a practical reason. Even if the axiomatic program is crowned with success and we are at long last in possession of rigorous full models of QFT, it is not likely that physicists will abandon EFT techniques or accord them a purely instrumentalist status. There is also a more conceptual reason that has to do with arguments we saw in chapter 2: RG techniques can also shed to the theoretical underpinning of rigourisation projects. Construed in a more local sense, the project of putting QFT on firm mathematical ground will need to consider each specific model separately to construct a full solution. It is not at all clear that the RG framework can be completely sidestepped or discarded. As a consequence, we have to agree with the more optimistic view that EFTs do indeed tell us something substantive about the world. But which realist-friendly stance is right?

We saw that the enthusiast overplays the hand they have been dealt. Besides issues of

internal coherence that have to do with the validity of the decoupling theorem and the form of the ontological picture drawn, there seems no external compelling reason to think that a future fundamental (final?) theory has been precluded. A full-hearted commitment to anti-foundationalism and anti-reductionism is premature. Additionally, save for evidence against the universal validity of the EFT program, so far we have found no evidence to indicate that we should exercise caution. Murphy's law notwithstanding, why not rejoice when everything is going according to plan? Hence, what we have labelled the optimist approach above seems to be winning the day.

Objection! ... Just as the jury was about to announce their decision in favour of the optimist camp, the devil's advocate revealed the last card up his sleeve: naturalness violation! New evidence will soon be presented, new witnesses will be called to the stand and the court will have to convene once more to reach their final verdict. For now, however, the session must be adjourned...

Chapter 3

When EFTs Seem to Fail I: The Hierarchy Problem

The standard model of physics, for all its unprecedented success in describing the world of microscopic phenomena, is expected to finally give way to a new theory that will account, among other things, for gravitational interactions. Apart from its descriptive incompleteness, indications of its eventual breakdown have been found in a series of unaddressed questions such as the preservation of the CP symmetry in strong interactions, the hierarchy problem in the weak sector and the cosmological constant problem when quantum fields are “allowed” to gravitate. These problems have been associated with violations of some “naturalness” principle which puts (meta)constraints on the kind of values the parameters of a theory can take. Various formulations of naturalness can and have indeed been given before the advent of the currently mainstream take known as “technical naturalness”. The standard response in the physics community was to take violations of naturalness as serious indications that there is something wrong with a given theory. Accordingly, attempted solutions to these problems have revolved around uncovering new physics in the form of supersymmetry, string theory or the multiverse in inflationary cosmology – while subscribing to the effective reading of QFT.

This is not the only possibility, however. A response, heretofore unpopular but slowly gathering some momentum in recent years, is to reject naturalness as a constraining principle and thus dismiss problems involving its violations as ill-formed or unmotivated. An arguably more radical alternative would be to consider the hierarchy and cosmological constant problems as indications, if not of a complete breakdown, of a potential limitation on the scope of validity of the EFT program. The important questions for this path would be to examine whether it

begets a thorough re-examination and re-evaluation of the way physics has been understood in later half of the 20th and early 21st century.

The structure of the chapter is the following. We begin with a brief motivating, semi-historical section aimed at presenting the most widely accepted definition of naturalness in the physics literature (technical naturalness). The specific way a problem is formulated often influences how much significance it is accorded. For this reason, section 2 presents three different formulations of the hierarchy problem in order of increasing persuasiveness with an eye on how each one corresponds to a violation of naturalness. Subsequently, three different classes of reactions are identified and sketched. The rest of the chapter comprises a careful critical examination of each reaction class. Section 3 begins with the standard (that is, prevailing among physicists) reaction to such problems: accept the they are meaningful and important and propose new physics to deal with or eliminate them. The status of naturalness as a guiding principle for theory construction is assessed. Section 4 turns to skepticism about the cogency and value of naturalness as an extra-empirical principle. Emphasis is given on fleshing out the assumptions of fine-tuning arguments and uncovering the physical insight – motivation behind naturalness. Section 5 entertains the possibility that the hierarchy problem signals a failure of the EFT program and assesses the repercussions this might have – especially with respect to reductionism. The emerging thesis is that unnatural theories fall outside the EFT paradigm scope when EFTs are understood in a “rigorous” as opposed to informal manner. Contra Wallace 2019, this does not have to be understood as rewriting the rule-book on how we do physics.

3.1 A Naturalness Story

Motivation The criterion or principle of naturalness has prominently featured in discussions of particle physics and beyond the standard model physics. Despite its close relation with these fields, the concept of naturalness can be expounded, to appropriate an example by Wells 2012 (§1.2) in the less exotic Galilean theory of falling bodies. As is well-known, Galileo was the first to fully realise that bodies left to freely fall from a fixed height on Earth move under the same acceleration g . It was this insight that drove him away from Aristotelian mechanics and its notion of an imparted force. Now, in analogy with work in particle physics, let’s assume, as is actually the case in retrospect, that Galileo’s theory is only an approximation to a deeper theory encompassing a wider range of phenomena. How would one go about modifying the

theory to take this fact into account?

First, one would try to find the kind of effects that could potentially alter the behaviour of the theory in some domains and incorporate them into the equations of motion. In this case, one might suspect that air resistance or the height at which a body is dropped might influence how fast the body will accelerate. In fact, we know that Galileo himself took air resistance to only have a negligible effect, so that it can be eliminated without sacrificing accuracy. His theory is already an approximation in this sense.

After deciding on the corrections necessary, one will include corresponding correction terms in the equation(s) of motion (or the Lagrangian (density) in modern physics) that will help fix the result to a more accurate value. Again, in the context of free fall, Wells adds corrective terms that account for changes in the value of g with respect to the height z such that the equation of motion

$$\ddot{z} = -g \quad (3.1)$$

changes into

$$\ddot{z} = -g + cz \quad (3.2)$$

with z obviously representing the height at which a body is released and c being a constant that will determine the significance of this correction. Given that z has dimension of length, dimensional analysis implies that c will have dimension of acceleration over length. Since the correction should not be comparable to the acceleration g , we can rewrite the constant as $c = \frac{g}{R}$ where R is of length dimension and lowers the value of the constant.

The key step now is to properly estimate the constant so that we can indeed produce more accurate predictions. In general, a good estimate for the constant will require some good estimate of the “right” scale of the problem. In the example we are examining, for instance, this comes down to finding a good estimate for R because this is the parameter that governs how the value of g is changing. Potential choices for R would thus involve the height of the highest point on Earth’s surface or that of the deepest point in the ocean or the Earth’s radius. The latter appears to be a more “natural” choice avoiding the extremities of either alternative. Although the range of possible values we can pick is admittedly wide (with differences up to orders of 10^3), we can include an additional parameter η to force the ratio closer to the right value, if

needed. Thus, the constant c ends up having the form:

$$c = \eta \frac{g}{R} \quad (3.3)$$

The naturalness requirement (at least according to one possible formulation goes – see below) would dictate that η be of order 1 or at the very least within a restricted range of, say, 10^{-2} to 10^2 (e.g. Giudice 2008, p. 8). Intuitively, this represents that our estimates have not fallen far off the actual values of the problem, i.e. the ratio $\frac{g}{R}$ only requires some minor adjustment to reproduce the right physics. Repeated measurements of bodies being released from different heights will lead to progressively more precise estimates for η .

Large Numbers The origins of naturalness are often (e.g. Giudice 2008, Giudice 2019, Craig 2017) traced back to the work of Eddington and a paper by Dirac 1937 on the constants of nature. Roughly, the idea, which has frequently been dubbed Dirac’s Large Numbers Hypothesis: “Any very large number occurring in nature should be simply related to a single very large number, which he chose to be the age of the universe” (Giudice 2008, §2). In effect, similarities between scale ratios should not be treated as mere coincidences but rather understood as insights towards uncovering deeper truths about the nature. Getting into the nitty-gritty of this discussion would only be a detraction here (for more see Giudice 2008). For our purposes, suffice to note that what fascinated people like Dirac in the early 20th century was the huge discrepancy of the proton mass compared to the Planck scale. A potential explanation for this came from Eddington’s intriguing observation that the number of protons N in the universe can be related to the gravitational force between a proton of mass m_P and electron of mass m_e :

$$\frac{e^2}{Gm_Pm_e} = \sqrt{N} \quad (3.4)$$

Dirac took this fact as a hint to a deeper fact about the universe, i.e. that the gravitational constant G would change with time. Why? Because this, according to him, was the only way for the left-hand side of the equation to remain equal to the constant on the right-hand side given the expansion of the universe over time. Thus, Dirac’s hypothesis can be seen as a guiding principle (or a motivating heuristic) towards a modified theory of gravity. It was soon noted, however, that this result (i.e. the evolving nature of G) would be incompatible with the formation of life on Earth. In what is probably the first true instance of anthropic

reasoning, Dicke 1961 estimated that the above ratio could be explained by contingent facts about the universe necessary to render the existence of observers possible. The awe-inspiring coincidence is exactly the value we would expect to measure by observers whose existence depended on time-scales large enough to allow for the formation of supernovae¹.

Although the connection between the above seemingly “kabbalistic” numerology and particle physics might appear rather loose, it is not hard to appreciate similarities between the two. First, note that in both contexts the explanandum is some vast scale discrepancy: the huge scale difference between the proton mass and the Planck scale in the cosmological coincidence above and the Higgs boson mass compared to the Planck scale for SM. What is more noteworthy, however, is the remarkable affinity in the style of reasoning involved each case. Recall:

- Large Number Hypothesis [LNH]: the mass ratios are such (natural?) that their combination with other parameters equals a large number which is some function of the age of the universe.
- Naturalness: the ratios between some parameters are the result of fine-tuning or cannot be explained through the violation of some symmetry (more about these soon).

In other words, the existence of these huge ratios [discrepancies] is acceptable because they reflect [they are due to] some deeper [at present perhaps unknown] fact about the structure of the universe. This signals a potential methodological device: when confronted with such bizarre numerical facts, be on the lookout for new physics! As we’ll shall soon, see while the Large Number Hypothesis does not have enough physical motivation, naturalness in high energy physics can be understood as possessing some more substantial content.

Before shifting gears, let us also underline this methodological analogy between Dirac’s LNH and naturalness in particle physics. As we’ve just mentioned, both principles are used to motivate the development of new physics: the former to motivate a modified theory of gravity and the latter, as we shall see, to motivate the introduction of new interactions such as those predicted by supersymmetry. But there is a deeper, more substantial commonality: they both point to specific issues (features) that the new theory will need to address (possess). In this way they can be more concrete guides for theory construction².

¹For more see on this story as well as criticisms of stronger versions of the so-called anthropic principle see Mosterin 2004.

²That is not to say that they can uniquely or unambiguously specify the path ahead. Dirac’s choice to treat

Definitions of Naturalness The first naturalness principle³, usually dubbed “Dirac Naturalness” codifies a numerological or “kabbalistic” abhorrence towards small numbers in Lagrangian expansions. For example, Craig 2017 states the principle as:

Dirac Naturalness: In a theory with a fundamental scale Λ , given an operator O of the form $L \supset c_O \times O$ with scaling dimension Δ_O , the natural size of the coefficient c_O in natural units is $c_O = O(1) \times \Lambda^{4-\Delta_O}$

This formulation precisely captures the idea that, when the scale dependence is made explicit, the parameters or any dimensionless constants should be of order 1. Obviously Dirac Naturalness can be too restrictive a criterion as certain parameters, which we would want to view as natural, could fail to pass the “of order 1” condition. If this is construed too strictly, for example, the fine structure constant, which is equal to $\frac{1}{137}$, of 10^{-2} , will also fail to be natural! This is why the requirement is frequently relaxed (e.g. Giudice 2008) and the range of natural parameters expanded to include values from 10^{-2} to 10^2 or even 10^{-3} to 10^3 . This amendment notwithstanding, the principle as such is still too restrictive. For example, as we will see in the following section, it will even rule out dimensional transmutation explanations for the proton-electron mass discrepancy.

The modern conception of naturalness began to emerge in the work of high energy physicists of the latter half of the 20th century. The driving idea was understanding the relation between low and high energy physics. In this vein, for example, Susskind emphasised the need to make a theory’s observable quantities largely insensitive to small variations of its fundamental parameters (Williams 2019b, p. 1031):

“Stability” Naturalness...a concept of naturalness which requires the observable properties of a theory to be stable against minute variations of the fundamental parameters (Susskind 1979, 2619)

the Cavendish constant as variable might seem arbitrary, but it is “a” project to pursue. The same applies, as we shall see with purported solutions to the hierarchy problem: there are multiple alternatives that can account for the apparent scale discrepancy. This is to be expected from an empiricist standpoint: the principles employed are not uniquely fixed by some metaphysical or rationalistic constraints or the form an old-fashioned Kantian would perhaps impose.

³There is a plethora of interesting papers on the historical development of conceptions of naturalness. Examples include Borrelli and Castellani 2019, Grinbaum 2012 and Williams 2019b but also the first few sections of Giudice 2008. The historical works nicely record how the concept developed through the early work of Weisskopf as an abhorrence towards quadratic divergences to later formulations emphasising parameters protected by symmetries. Here we simply survey some key stops in this development to arrive at the most widely used formulation of the principle.

This principle ensures that one can remain largely ignorant of the exact details of the microscopic regime and still make accurate predictions about the larger length scales they can access. The problem, according to this view, with quadratic divergences of the form we find in the case of the Higgs boson mass is that they clearly, as we shall soon see, instantiate a form of undesirable UV sensitivity: heavier masses contributing to quantum corrections fail to decouple (Schwartz 2014, p. 408-10).

While this might initially appear to be the case for all particles –after all we find loop diagrams and terms proportional to large scales for all of them when we attempt to renormalize–, symmetry considerations single out scalar fields as particularly vulnerable to this sensitivity. Indeed, fermions or gauge bosons will be “protected” from such sensitive dependence thanks to symmetries like chirality and gauge invariance. This motivates a further suggestion for a stricter definition of naturalness by ‘t Hooft, technical naturalness, which accords centre stage to symmetries:

‘t Hooft Naturalness: “at any energy scale μ , a physical parameter or set of physical parameters $a_i(\mu)$ is allowed to be very small only if the replacements of $a_i(\mu)$ would increase the symmetry of the system” (‘t Hooft 1980, 136)

The recent reformulation of this technical version of naturalness, which is applied in BSM physics can now be stated, following again Craig 2017, as:

Technical Naturalness: Coefficients can be much smaller than their Dirac natural value if there is an enhanced symmetry of the theory when the coefficient is taken to zero. In this case, the natural size of the coefficient c_O is $c_O = S \times O(1) \times \Lambda^{4-\Delta_O}$ where S is a parameter that violates the symmetry in question.

Effectively, this means that very small dimensionless parameters may be present in the perturbative expansion on condition that, when set to zero, they make manifest some symmetry lurking beneath the effective description. For example, a very small fermion mass term is tolerable if setting it to zero “enhances” the symmetry of the Lagrangian - which in this case is true: the Lagrangian becomes symmetric under chiral transformations. The (physical) reason for the “protected by symmetry” clause is that when corrections are associated with a symmetry, they will need to enter calculations as terms proportional to the couplings of those symmetry breaking terms (think of the way counterterms are added to the Lagrangian – explicitly mimicking the structure of, or being proportional to, the terms they “counter”). This implies that

the smaller the couplings associated with these terms, the more heavily the symmetry violating corrections will be suppressed (since they will be proportional to some very small quantity).

It is not hard to appreciate how this “technical” version of naturalness can act as a guiding principle to developing new theories: At first, an unnatural parameter will be perceived as an indication that there is something unsatisfactory about the description currently available. The way towards a more satisfactory description, implicating new physics, will be to postulate a symmetry that is violated (or “broken”) when a term with the corresponding unnatural parameter is included in the Lagrangian ⁴. When the symmetry breaking term is taken to zero the (postulated) symmetry should be restored. This is how approaches like technicolor or supersymmetry, which expand the Lagrangian symmetries and introduce new degrees of freedom, are motivated as “natural” solutions to the hierarchy problem.

3.2 The Hierarchy Problem

There are many indications that the standard model is not the final word in fundamental physics: apart from its obvious failure to incorporate gravitational interactions, one can point to the non-zero neutrino masses, the failure to unify all interactions under one common group or perhaps more dramatically the existence of a Landau pole for QED. In the language of EFTs we anticipate a breakdown, a failure of the theory to be predictive, that is, to produce meaningful results at that scale. It has long been anticipated that the SM will collapse well before the QED Landau pole is reached, however. Although the possibility that this cut-off is closer to the presently accessible scales has not been ruled out, the most widely accepted scenario is that this cut-off scale will be the Planck scale (around 10^{19} GeV), a scale at which contributions from the gravitational interaction will cease to be negligible. Here is a first puzzle. As we know, the electroweak scale, the energy scale \mathcal{V} at which most processes described by electroweak theory occur is around 246 GeV. Whence this huge discrepancy⁵?

$$\frac{\mathcal{V}}{M_P} \approx 10^{-17} \quad (3.5)$$

We already examined Dirac’s puzzlement over the discrepancy between the mass of the electron compared to that of the proton. The hierarchy problem essentially is a similar puzzle-

⁴Again, think of chiral symmetries: the mass terms $m \bar{\psi}\psi$ of the fermions can be really small because their appearance breaks the chiral symmetry of the corresponding massless Lagrangian

⁵See figure 3.1 taken from Craig 2017 (p. 15).

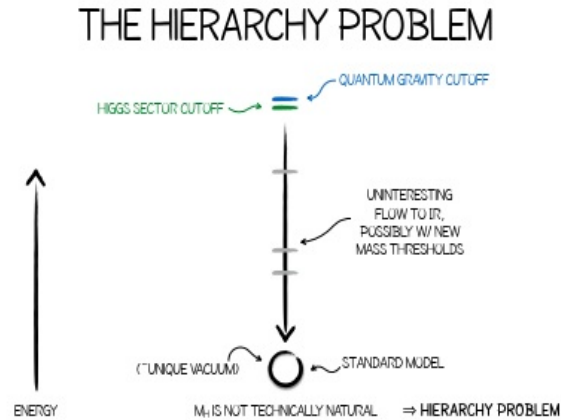


Figure 3.1: Schematic Representation of the Hierarchy Problem

ment about the discrepancy between the Planck scale and that of the electroweak scale. It is frequently put in the form of the following question:

Hierarchy P1: why is the gravitational interaction so much weaker (order of magnitude $\approx 10^{-25}$) than the electroweak interaction?

Prima facie, questions of this sort seem moot. After all, one does not have to be an extreme empiricist to accept that at least some of our knowledge will involve brute facts, which we will be unable to derive or justify using other (empirical) facts or, for those not allergic to rationalistic enterprises, through appeal to some transcendental-like (metaphysical) principles. The hierarchy problem thusly presented might be reminiscent of an obstinate child’s refusal to halt an endless sequence of “why” questions. Alas, explanations have to end somewhere, as Wittgenstein would say. Nonetheless, physicists do find such questions worth pursuing as potentially conducive to important conceptual breakthroughs. We will evaluate this claim in the following section when we examine allegedly successful instances of applying naturalness.

First, let’s take a look at the hierarchy problem through a more technical lens. From a mathematical perspective, the whole issue is essentially a problem of “false expectations” involving the mass of the Higgs boson. There are different ways to present the problem and each one highlights slightly different aspects of its impact –a fact that we will exploit when approaching the various reactions to it–, but sufficiently rudimentary way to present it is the following. As we saw when examining the process of renormalization, infinities in the various Feynman diagrams containing loops were essentially “absorbed” in the parameters of the theory. This means that these parameters were carefully adjusted to generate terms which would eliminate the infinite parts in the expansions and leave as a residue a physically significant finite part.

We tried to alleviate worries about this cancellation of infinities by “infinite” parameters arguing that, after all, bare parameters only indirectly derive their physical meaning through the renormalisation procedure.

Infinities in loop diagrams for scalar fields are tamed in the same manner. Drawing on Schwartz 2014 (p. 408-9)⁶, we present an example for a scalar field of mass m coupled to a fermion field of mass M . In this case we need to insert fermion loops as corrections to the propagation of the scalar field (diagram taken from Schwartz 2014):




Figure 3.2: Fermionic Loop in Scalar Field Propagator

These loop corrections give rise to divergences that need to be treated with a renormalization procedure for the scalar fields and their mass parameter. To understand the connection with the above characterization of the hierarchy problem we must start with a Lagrangian describing the coupling between a scalar field ϕ of mass m and a fermion field M (with $m < M$):

$$\mathcal{L} = -\frac{1}{2}\phi(\partial^2 + m^2)\phi + \bar{\psi}(i\not{\partial} - M)\psi + \lambda\phi\bar{\psi}\psi \quad (3.6)$$

we note that the contribution from the above diagrams is:

$$i\Sigma(\not{p}) = -(i\lambda)^2 \int \frac{d^4k}{(4\pi)^4} \frac{\text{Tr}[(\not{p} + \not{k} + M)(\not{k} + M)]}{[(p+k)^2 - M^2 + i\epsilon][k^2 - M^2 + i\epsilon]} \quad (3.7)$$

which, using e.g. dimensional regularisation will give:

$$\Sigma(p) = \frac{-\lambda^2}{4\pi^2} \left[\frac{6M^2}{\epsilon} - \frac{p^2}{\epsilon} + M^2 - \frac{p^2}{6} + \int_0^1 dx [3p^2x(1-x) - 3M^2] \ln \left(\frac{M^2 - p^2x(1-x)}{4\pi\mu^2 e^{-\gamma}} \right) \right] \quad (3.8)$$

with m the scalar field mass, M the fermion field mass. \not{p} the external momenta for the scalar field, k the internal momenta for the fermion field and λ the coupling constant.

Treating the SM as an EFT of a more fundamental theory, we know that the Higgs mass appearing at the electroweak scale, the scale at which the symmetry of the Lagrangian is broken, must be related to the corresponding mass at the (more “symmetric”) high-energy Planck scale.

⁶See also section 2.1 of Craig 2017

Since other fields such as the fermion field appearing as part of the quantum corrections have are heavier, i.e. their mass M is much higher than that of the Higgs field m , it follows that the mass of the Higgs boson should also be very large in complete disagreement with the measured value. We can thus put the problem in a slightly less informal way:

Hierarchy P2: Why is the Higgs boson mass m_{EW} so much smaller at the electroweak compared to its value m_Λ at the Planck (or whatever more fundamental, high energy) scale Λ ?

We have now established the connection between the hierarchy problem as a puzzlement over the scale discrepancies of fundamental interactions and as a problem with small empirical value for the Higgs boson mass. Next, we will transcribe the problem in a more technical language by emphasising what has been taken to be its most disconcerting feature: fine-tuning.

3.2.1 Formulations

We offer three possible formulations of the problem, each emphasising some distinct (but interconnected) aspect of the problem. The formulations are given in an increasing order of persuasiveness, at least with respect to the “stakes” involved in each case.

FORMULATION I: fine-tuned bare mass One straightforward way of presenting the problem is the following. If we assume that a more fundamental theory T exists which describes the world at a (more) fundamental scale Λ (e.g. in string theory the length of strings, the Planck scale 10^{19} GeV), then the “bare” parameters in the Lagrangian would also be physically significant - corresponding to the values of these parameters at this more fundamental case. In this case, for the pole mass of the Higgs boson m_P we would have:

$$m_P^2 = m_0^2 + \delta m^2 \tag{3.9}$$

with m_0 now being the physical (bare) mass and δm representing the higher-order quantum corrections coming from fields interacting with the Higgs field (e.g. quarks such as the top quark whose coupling is y_t , the massive gauge bosons such as gluons and Z^0, W^+, W^- with couplings g, g' and so on):.

$$\delta m^2 = \frac{\Lambda^2}{16\pi^2} \left(-6y_t^2 + \frac{9}{4}g^2 + \frac{3}{4}g'^2 + 6\lambda + \dots \right) \tag{3.10}$$

The problem arises when we finally decide to take off our theorist’s hat and put on the experimentalist’s in order to obtain some empirical estimates for the mass value at a particular energy scale (currently around 10^3 GeV at the LHC). The story then goes as follows. First, we note:

- Calculations for the corrections indicate a quadratic dependence on the (fundamental) scale $\delta m^2 \sim \Lambda^2$
- The value of the scale is expected to be around the Planck scale $\Lambda \sim 10^{19}$ GeV or at the very least some value above the currently accessible scales at the LHC $\Lambda \sim 10^3$ GeV
- Recent experimental results at the LHC indicate that physically measured mass is around $m_P \approx 125$ GeV, i.e. of order 10^2 .

To make the theoretical estimate of the mass agree with the measured value of the physical mass, the bare parameter would need to be carefully adjusted to cancel the huge quantum corrections. In other words, we would need the bare mass to be scale with Λ like:

$$m_0^2 = m_P^2 - \delta m^2 \approx \epsilon \Lambda^2 + \Lambda^2 = (1 + \epsilon) \Lambda^2 \quad (3.11)$$

Clearly, ϵ needs to be of the right order to bring the scale of the renormalised mass down to order 10^4 of m_P^2 . Given that Λ^2 is of order $(10^{19})^2$ it follows that $m_0^2 \approx (1 + 10^{-34}) \Lambda^2$. In other words, the bare parameter would need to be adjusted to 33 decimal digits to reproduce the physically observed value for the Higgs mass. Even if the fundamental scale is lower than 10^{19} GeV (perhaps even as low as only slightly above the energies most recently triggered at the LHC), the bare parameter would still need to match the corrections up to several decimals, as determined by the ϵ coefficient.

Whether or not the parameters are treated as physical, it appears too fortunate a coincidence for the bare mass to have the ridiculously specific value needed to exactly cancel the quantum corrections and reproduce the physical value at low energies. In other words, the bare parameter needs to be **fine-tuned** or precisely chosen in order for the theory to produce accurate results. This fact is typically seen as “crying out for an explanation” and has thus attracted considerable attention. To be sure, rendering this notion of “coincidence” more accurate will be one of the challenges for substantiating this aversion to “fine-tuning”.

FORMULATION II: delicate matching condition The above characterisation of the problem in terms of bare parameters of a more fundamental theory will probably seem unconvincing

to someone that is inclined to assert that bare parameters should be treated as formal tools for producing correct calculations and only agrees to confer physical significance only to renormalised parameters which alone suffice for evaluating the correlation functions. Wetterich 1984 (p. 217), for example, after distinguishing between two types of fine-tuning problems where “the first fine-tuning problem concerns the relation between physical quantities (like the W boson mass) and the bare parameters appearing in the action used to define the functional integral describing the model”, nicely captures this sentiment:

In any case, fine-tuning of bare parameters is not really the relevant problem: we do not need to know the exact formal relation between physical and bare parameters (which furthermore depends on the regularization scheme), and it is not important if some particular expansion method needs fine-tuning in the bare parameters or not. The relevant parameters are the physical parameters, since any predictions of a model must finally be expressed in terms of these.

Rosaler and Harlander 2019 go even further treating the scale Λ in Wilsonian renormalization as a formal tool that is chosen arbitrarily for the sake of extracting physical results from a QFT. They even go as far as to recommend a whole new way of conceiving of EFTs, namely as identifiable with full trajectories in parameter space:

In this way of understanding Wilsonian EFTs, an EFT is specified by an entire Wilsonian RG trajectory rather than by any single point on such a trajectory, and points along the RG trajectory are understood as physically equivalent parametrizations of the same set of physical quantities.

This possibility notwithstanding, we know that there is an alternative, “continuum” version of the renormalization group. It is thus important to see whether the problem can be reformulated in a way that will not depend on special features of the Wilsonian approach and implicate the renormalized instead of the bare parameters. In this scenario, one will need to phrase the problem in terms of **matching conditions** between the parameters of the standard model as an EFT and a more complete, underlying theory T .

Recall our standard procedure in this case: we start from the action S_{full} for the more fundamental, full theory, which we evolve to lower energies $\mu' = \mu - \delta\mu$ using the renormalization group equation. At a certain scale $\mu = M$ of interest we can replace the full action with an effective action S_{eff} , which only contains fields lighter than M . Its coefficients are evaluated through the matching procedure to guarantee agreement between the two theories. When

turning to the mass of a light scalar field like the Higgs boson field, we would have

$$m_h^2(\mu) = m_H^2(\mu) - \delta m_H^2(\mu) \quad (3.12)$$

with m_h the Higgs mass in the SM, m_H the Higgs mass in the more encompassing theory and δm_H the corrections – contributions to make these two parameters agree. It is important to note that all parameters involved now are renormalized, running parameters and not the bare kind we saw above.

Harlander and Rosaler 2019 nicely illustrate the problem by revisiting a standard toy example (e.g. Burgess 2007) of deriving the effective Lagrangian of a light field ϕ from the Lagrangian of the light scalar field coupled to a heavy fermion ψ . The effective Lagrangian will have the following form:

$$\mathcal{L}_{eff} = \frac{1}{2}A(\partial_\mu\phi)^2 - \frac{1}{2}B^2 + \frac{1}{4!}C\phi^4 + \dots \quad (3.13)$$

with the coefficients A, B, C left to be specified through the matching procedure (note how higher order terms are here neglected as they will not be relevant at the lower energies examined). Recall that the matching condition can be extracted by comparing the propagators G_{2eff} and G_2 of the two theories, which contain contributions from the fermion field, and will lead to a relation between the masses of the scalar field in the two theories of the following form:

$$m_h^2(\mu) = m_H^2(\mu) - f(M^2) \quad (3.14)$$

with m_h the scalar mass at the effective theory, m_H at the full theory and f a function of the thresholds mass scale M of the fermion field corresponding to its contributing corrections. For the sake of brevity we do not give a specific form to this f function⁷. In this context, the hierarchy problem corresponds to the fact that only a very limited range of values for μ will realise the necessary delicate cancellation between the full theory mass and its corrections to produce the small value of the mass in the effective theory - which is, after all, the value we measured in recent experiments.

Although ignoring the problem as a mere technicality concerning unphysical bare parameters is no longer an option in this case, detractors of naturalness might still insist that this version

⁷See also Schwartz 2014 (p. 409) for detailed calculations and results using two different renormalisation schemes.

of the problem, pending a more convincing case than the mere undesirability of fine-tuning, fails to impress violations of naturalness as disconcerting facts. If this reply feels somewhat desperate at this point (given that we are talking about physically relevant parameters now), the sceptic has a few more weapons at their disposal, as we shall see in section 3.4. The main theme will be to argue that the problem is an artifact of the regularisation or renormalisation schemes chosen.

FORMULATION III: defective operator behaviour Perhaps the most forceful formulation of the problem is the one which most acutely stresses its inimical implications for the EFT program undermining the very viability of an effective approach to QFT. This is done by revisiting the Lagrangian expansion for an EFT and examining the behaviour of its coefficients as the energy scales μ are altered. To appreciate this point, let us recall that an EFT is essentially an infinite sum of operators of the form:

$$\mathcal{L}_{eff} = \mathcal{L}_n + \mathcal{L}_{n+1} + \mathcal{L}_{n+2} + \dots \quad (3.15)$$

where each \mathcal{L}_n contains operators with dimension m , for all $m > n$ and n being the spacetime dimension. These operators, tracking the interactions at increasing energies, scale like (see Petrov and Blechman 2015 p. 57):

$$\mathcal{L}_{n+m} = \sum c_i O_i \sim \left(\frac{p}{\Lambda}\right)^m \quad (3.16)$$

with their coefficients c_i determining the likelihood of the corresponding processes⁸. As is easily seen, the values of the coefficients change with the energy scale for $m > 0$. Dimensional analysis can be used to extract their behaviour as a function of the energy scale. The corresponding operators can be classified as: i) relevant, ii) irrelevant and iii) marginal operators. Recall that irrelevant operators [$m > 0$] become less and less important as the energy scale becomes lower (or equivalently the length scale becomes higher), relevant operators [$m < 0$] exhibit the opposite behaviour, that is, their significance increases with the decrease in energy (or increase in length scale) while marginal operators [$m = 0$] have the same contribution across all levels and their behaviour requires a more careful analysis. Generally speaking, relevant operators will significantly interfere with low-energy physics and tracking their behaviour

⁸For example, one such term could be $\lambda^{(6)}\phi^6$ in the simple case of scalar theory.

will be essential if one is to extract accurate results.

In the standard model, the Higgs boson mass is the only coupling constant of a relevant operator (much like the mass term in ϕ^4 theory)

$$\mathcal{L}_2 = \frac{m_H^2}{2} H^\dagger H \quad (3.17)$$

Thus, we expect that, as the energy scale is getting lower (or the length scales are increased) the Higgs mass term will become increasingly higher. Equivalently, we expect that as we use the RG transformations to go to greater length scales, the Higgs mass value will become very big as well. In fact, this is a massive understatement. For since the mass scales like Λ^2 : $m_H^2 \sim \Lambda^2$, its value (even on the modest of estimates) for the Λ scale will be colossal! This is the contradiction with our observations which tell us that the Higgs mass at long distances is a remarkably smaller number [125 GeV] than what our analysis would lead us to believe. Clearly, something does not add up.

This formulation of the problem mostly vividly illustrates the way violations of naturalness can undermine the RG tools we use along as part of the effective program in QFT. Indeed, using a toy example found in section 23.6.1 of Schwartz 2014⁹, we can appreciate the way in which the RG equations need to be fine-tuned in order to deliver the correct values for relevant operators. Let us consider a theory with dimension 4 and dimension 6 operators, whose corresponding couplings are labelled g_4 and g_6 . For convenience we redefine the constants using: $\lambda_4 = g_4$ and $\lambda_6 = \Lambda^2 g_6$. These new coupling constants now evolve as:

$$\Lambda \frac{d}{d\Lambda} \lambda_4 = \beta_4(\lambda_4, \lambda_6) \quad (3.18)$$

$$\Lambda \frac{d}{d\Lambda} \lambda_6 - 2\lambda_6 = \beta_6(\lambda_4, \lambda_6) \quad (3.19)$$

When solving these equations for small β -s we use a linear approximation with parameters a, b, c, d :

$$\Lambda \frac{d}{d\Lambda} \lambda_4 = a\lambda_4 + b\lambda_6 \quad (3.20)$$

$$\Lambda \frac{d}{d\Lambda} \lambda_6 = c\lambda_4 + (2 + d)\lambda_6 \quad (3.21)$$

We then diagonalise the system so that solutions are easy to obtain and then return to the

⁹A very similar example is also examined by Williams 2015.

original basis and parameters to impose some boundary conditions. For simplicity here, one might take $\lambda_4(\Lambda_H) = \lambda_6(\Lambda_H) = 0$ and arrive at the following:

$$\lambda_6(\Lambda_L) = -\frac{c}{2} \left(1 - \frac{\Lambda_L^2}{\Lambda_H^2}\right) \lambda_4(\Lambda_L) \quad (3.22)$$

$$\lambda_4(\Lambda_L) = \frac{b}{2} \left(1 - \frac{\Lambda_H^2}{\Lambda_L^2}\right) \lambda_6(\Lambda_L) \quad (3.23)$$

with Λ_L some low energy scale and Λ_H a much higher one: $\Lambda_L \ll \Lambda_H$. Treating the higher energy value as an initial condition, we want to see how sensitively the values of the couplings at the lower energy depend on those at higher scales. As Λ_H becomes larger, the limit for λ_6 converges to a finite value. This coupling constant is irrelevant: the high energy value does not influence its value at lower scales.

Things are dramatically different for λ_4 , however. As $\lambda_H \rightarrow +\infty$, $\lambda_4(\Lambda_L)$ fails to converge to a finite limit. The low energy value of λ_4 sensitively depends on its value at high energies when λ_6 is held fixed. The Higgs mass term is completely analogous to this λ_4 coupling constant. In order to fix the low energy value $m_H(\Lambda_L)$ at the experimentally derived value, one would need to give a very precise value to the high energy $m_H(\Lambda_H)$:

$$m_H(\Lambda_H)^2 = \frac{m_H(\Lambda_H)^2}{\Lambda_H^2} \sim 10^{-34} \quad (3.24)$$

Tiny changes of the high energy value of the mass will lead to incredibly huge changes at the low energy values. It appears that for the theory to get things right at the scales we can access experimentally, nature has to “conspire” at (more) fundamental scales to fix the higher energy constant within the very narrow range of values that will be compatible with the low energy result we observe. So much for the “decoupling of scales”!

Common Diagnosis : Note how all of the above cases constitute examples of naturalness violation according to each definition of naturalness we saw above:

- **Dirac**: it is straightforward to see that the Higgs mass fails to satisfy this constraint given the remarkably small value that needs to be attributed to the bare or high energy value to recover the correct empirical value.
- **Technical**: there is no symmetry protecting the Higgs mass to make the noted discrepancy between scales and the tiny value attributed to the bare parameter legitimate.

- **Autonomy**: note how the need to fine-tune the high energy value in the third formulation seems to imply that we need knowledge of the low energy scales to fix those at high energy scales.

It is about time we saw what led us astray and what can be done to resolve or dissolve this problem in BSM.

3.2.2 Reactions

Now that we have familiarised ourselves with some of the historical motivation and definitions of naturalness as well as varied formulations of the hierarchy problem, it is about time we gave a brief overview of the possible stances one can take on it. The rest of the chapter will be devoted to a systematic examination of the reasoning and philosophical implications of the key reactions to the problem. We will find that there are three main axes along which to continue:

1. New Physics : The currently prevalent position within the scientific community is to accept the significance of the problem and attempt to treat it by developing new BSM physics (e.g. Craig 2017, Dine 2015, Feng 2013). Multiple extensions or revisions of the standard model have been proposed that could potentially eliminate the quadratic divergence e.g. by introducing additional symmetries which restore technical naturalness (supersymmetry) or by seeing the Higgs boson as a composite non-bosonic particle (technicolor), adding dimensions (e.g. in string theory), or through anthropic reasoning in the multiverse (e.g. Susskind 2003, Susskind 2008, Weinberg 2007). In all cases, one agrees that naturalness is a well-formulated and physically meaningful principle that acts as a constraint on possible future physical theories (though its exact character varies significantly in the context of anthropic reasoning – Borrelli and Castellani 2019, Grinbaum 2012, Williams 2019b). The main obstacle in this path is... experience: recent runs of the LHC at TeV energy scales have furnished no evidence for any of these alternative theories (Craig 2014, Giudice 2019). As a result, an important philosophical question arises about the nature of naturalness: is it a heuristic or a “physical” principle¹⁰?) and, more broadly, the role of heuristic principles in guiding theory construction. As is known, the advent of the EFT program led to an abandonment of renormalizability as a constraint on their

¹⁰This distinction is meant to capture the difference between principles that used during the stage of theory development but for the most part mark a contingent or non-physical feature of a theory and those that do capture (at least) part of its physical content. An example of the former would be general covariance in GR, whose status as a physical principle becomes largely deflated under scrutiny. (see Norton 1993). An example of the latter would be Einstein’s postulate about the constancy of the speed of light with respect to inertial frames in SR.

predictive status of theories (as expressed e.g. by Huggett and Weingard 1995) and saw the rise of naturalness. Given these recent misgivings, can naturalness be admitted as a principle that has historically led to novel developments or should it be seen more as part of a reconstructive project (post-rationalisation) of this history (Wells 2012, Wells 2015, Sahuquillo 2019)?

2. Denial : Sceptical voices seem to be gaining some ground in recent years with philosophers (Rosaler and Harlander 2019, Bain 2019) and physicists alike (Hossenfelder 2018 but also Manohar 2018 or Wells 2019) questioning the significance of the problem and the validity of naturalness as a (guiding) principle. Perhaps the whole problematic, these sceptics argue, is a consequence of poorly conceived premises that collapse under closer scrutiny. There is definitely something true about this suspicion: frequently, the persuasiveness naturalness-oriented arguments heavily depends on the exact formulation of the problem - because this is what determines the “price” one has to pay for disregarding it. For instance, if the demand for naturalness is presented in probabilistic terms (e.g. à la Hossenfelder 2019), that is as underlining the incredibly small chance that fundamental parameters will take on the values for which the theory is rendered predictive, then rejecting the argument as a non-scientific numerological pseudo-problem will probably seem innocuous. As part of this sceptical onslaught, we will present three escalating “attack waves” against the cogency of the problem:

- A Technical objections raised by those reducing the problem to a poor choice of regularisation scheme (e.g. Manohar 2018) or dissolving it by an appropriate choice of renormalisation scheme: on-shell renormalisation (Bain 2019, Harlander and Rosaler 2019). The main problem with these claims is that they disregard the dependence of lower energy physics to higher energy particles that should drop out of the picture at low energies.
- B Disputing the status of fine-tuning by arguing that there is no well-posed problem because no well-motivated or non-question begging probability measure is given (Hossenfelder 2019). Without a properly motivated distribution no conclusions about the likelihood or unlikelihood of a theory can be reached (Wells 2019).
- C Contesting the motivation, i.e. the physical insight, behind naturalness (e.g. Hossenfelder 2018). Examining some candidates (e.g. those found in Williams 2015) like aversion to quadratic divergences, fine-tuning and beauty (Donoghue 2017) and how it should be understood – with simplicity being an obvious candidate –, we will turn to what Williams

labels the “autonomy of scales” as the strongest and most physically meaningful content for naturalness.

3. EFT Failure : Perhaps the violation of naturalness and the problems with the Higgs boson mass are a clear indication that the decoupling of scales we assumed in the Wilsonian approach to renormalization was wrongheaded from the very beginning. Perhaps nature is not so kind to us after all and what we need to do is come to terms with the “disconcerting” fact that input from macroscopic physics somehow is indispensable in understanding the microscopic regime. This means that, similar to chaotic phenomena, details about specific micro-arrangements are not washed out under the dynamical evolution of the system, but actually need to be carefully specified to make sense of its current state. What would the repercussions be when we abandon this commitment to naturalness as a decoupling of scales? Wallace 2019 has equated the rejection of naturalness with a rejection of the reductionist program that has heretofore been so successfully applied in science. We will dispel this worry by, first, disentangling naturalness from decoupling. Invoking recent philosophical work on EFTs by Franklin 2020 and Bain 2019, we will try to ground the success of EFTs in a less narrow conception of decoupling than that designated by naturalness. Lest this be seen as unconvincing, we will further undermine the scope of Wallace’s conclusion by highlighting the disanalogies between QFT and other fields of physics like statistical mechanics by drawing on relevant work by D. Fraser 2020 and D. Fraser and Koberinski 2016 on the Higgs mechanism and renormalization group techniques in both frameworks.

3.3 Reaction 1: New Physics*

Extensions and modifications of the SM have been presented for various reasons, but naturalness has certainly been an additional or even the main motivation for some of these¹¹. Some of these solutions include (for more technical details one could consult sources such as Craig 2017, Giudice 2019, Hebecker 2021, Koren 2020, ch. 3):

1. Lower Cut-Off An obvious way to render the fine-tuning less pressing a concern is to lower the cut-off energy scale. This move potentially increases the range of possible values

¹¹Readers in a hurry can skim through the first half of this section without remorse. They can focus on the more philosophically interesting section 8.3.1 which analyses the historical role played by naturalness in the development of the SM.




Figure 3.3: Compactification: the additional dimension of the superspace is “winded” around the subspace of the effective theory

for the bare mass to render the fine-tuning more palatable. After all, it’s one thing to be hit by a meteorite and another to be struck by lightning! One way to achieve this “deflation of unlikelihood” is to postulate extra spatiotemporal dimensions. How does this have anything to do with the value of parameters? Essentially, it all comes down to the determination of dimension of the coupling constants via dimensional analysis. To illustrate this point, consider Newton’s gravitational constant (Dine 2015, p. 10). Start with the following Lagrangian:

$$\mathcal{L}_{d+4} = \kappa_{d+4}^{-2} \int d^4x d^d y \sqrt{g} R \quad (3.25)$$

with d indicating the number of extra dimensions and κ the modified Newton constant. The addition of d changes the power counting estimate so that the constant scales as $\kappa^2 \sim (TeV)^{-(2+d)}$. Clearly, as the number of dimensions goes up, the value of the constant drops. The Higgs boson in the standard model may be “cured” in an analogous manner with the introduction of objects called 3-branes, whose excitations correspond to particles in 4 dimensions. In string theory these branes “live” in some space of dimensionality greater than that of ordinary space-time. The additional unobserved dimensions are hidden from view through a process known as compactification. The compactified dimensions lead to an increase in the effective Planck mass compared to the fundamental Planck mass resolving the huge discrepancy problem¹².

2. Technicolor Having noted that scalar fields are peculiar in their lacking some symmetry “protection”, a reasonable move would be to simply exclude them from a theory. Technicolor is the mathematical realisation of this insight by treating the Higgs boson as a composite, rather than a single, field. Technicolor essentially offers an alternative mechanism for electroweak

¹²More technical details on the mechanism of Kaluza-Klein reduction and other strategies of similar kin like orbifold reduction – see Koren 2020 sect. 3.2.

symmetry breaking that does away with the complex Higgs field. Early models in this program essentially mimicked the symmetry breaking mechanism we find in the formation of mesons by QCD quarks. Adding a new gauge interaction $SU(N_{TC})$, called “techni-color”, which, similarly to the strong interaction, is asymptotically free at large energies but becomes stronger at larger distances. The phenomenology derived from this new gauge group is completely analogous to the that for the strong force (e.g. one derives formation like “technipions”). One, then, embeds this group to a larger group $SU(N_{TC}) \times SU(3)$ to make sure that the various fermions couple to the quarks (so as to be obtain masses). Breaking this group’s symmetry through SSB results in a similar physical situation to that obtained through the Higgs mechanism.

3. Supersymmetry The clearest way technical naturalness can serve as a guide to new physics is by mimicking the strategy we saw at work in the case of fermions: take the unnatural parameter to indicate the presence of an underlying symmetry. The most popular approach in this camp, of course, is supersymmetry (SUSY). SUSY generalises the symmetries of the SM by introducing so-called superpartners for all known fields. At its core, this amounts to SUSY postulating a relation between internal and external symmetries¹³ so that all fermions and bosons now acquire their respective bosonic and fermionic superpartners. The violation of naturalness is thus explained through the breaking of this symmetry - much like a fermion mass breaks chiral symmetry. The end result is that these new superfields introduce contributions in the correcting loops that cancel the problematic quadratic terms issuing from the SM fields.

Blitz SUSY Primer It is worth exploring how SUSY succeeds in eliminating these divergences in some more detail. To achieve this, however, we will need some background first (Duncan 2012, sect. 12.6). Recall that symmetries in the SM are connected to Lie Groups: the Lorentz group for boosts and rotations in spacetime and the Poincaré group when spacetime translations are included as well. For internal symmetries the groups are $SU(3)$, $SU(2)$, $U(1)$, with $SU(3)$ describing the strong and the $SU(2) \times U(1)$ the electroweak interaction. The groups can be represented (among other things) by matrices acting on vectors, spinors etc. Their action is produced by operators of the form e^{iJ} , with J being the generators of the group (starting from a generator we can produce the full operation through multiple iterations of infinitesimal

¹³This means that transformations for internal symmetries for something like the spin degrees of freedom and to transformations of a spatiotemporal kind essentially constitute manifestations of a more fundamental unified symmetry group – much like spatial and temporal transformations were found to be part of a more fundamental group in relativity theory.

transformations δJ). These generators form Lie Algebras defined by Lie brackets. For external symmetries the Lie brackets are given by:

$$[P_\mu, P_\nu] = 0 \quad (3.26)$$

$$[M_{\mu\nu}, P_\rho] = i(g_{\rho\nu}P_\mu - g_{\rho\mu}P_\nu) \quad (3.27)$$

$$[M_{\mu\nu}, M_{\rho\sigma}] = i(g_{\rho\mu}M_{\rho\nu} + g_{\rho\nu}M_{\mu\rho} - g_{\rho\mu}M_{\sigma\nu} - g_{\rho\nu}M_{\mu\sigma}) \quad (3.28)$$

The Lie algebras for internal symmetries are of the form:

$$[g^\alpha, g^\beta] = i f^{abc} g^c \quad (3.29)$$

$$[g'^i, g'^j] = i \epsilon^{ijk} g'^k \quad (3.30)$$

$$[g_0, g_0] = 0 \quad (3.31)$$

with g, g', g_0 representing the 8, 3, 1 generators for the gauge groups of the strong, weak and electromagnetic interactions respectively.

Note that the above framework prohibits any transformation between fermion and boson fields: there are no Lie brackets such that $[L, g] = 0$, where $L = M_{\mu\nu}, P_\mu$. This is a direct consequence of the **Coleman & Mandula theorem**, which states that, under some fairly innocuous assumptions¹⁴, the Poincaré group is the most general symmetry group we can find (modulo direct products with itself). The key idea, and starting point for SUSY, is to evade this theorem by dropping one seemingly plausible assumption: that the generators of the symmetry transformations can only satisfy commutation (and not anti-commutation) relations. This is equivalent to denying that the generators of the symmetry group need be bosonic. By changing the spin of a state by $\frac{1}{2}$, SUSY operators can transform bosons into fermions and conversely

$$[SUSY] |boson\rangle = |fermion\rangle \quad (3.32)$$

$$[SUSY] |fermion\rangle = |boson\rangle \quad (3.33)$$

To do so, one first extends the Poincaré group in the simplest possible way by introducing

¹⁴For example, that all particles correspond to representations of positive energy and for a given m the number of particles with mass below m is finite.

spinorial generators of the form:

$$Q_a = \begin{pmatrix} C_s Q^* \\ Q \end{pmatrix} \quad (3.34)$$

with $\alpha = 1, 2, 3, 4$, Q a $(0, \frac{1}{2})$ 2-spinor and C_s the conjugation matrix. The generators obey anti-commutation relations with one another, but still commute with the momentum generators:

$$\{Q_\alpha, \bar{Q}_\beta\} = 2(\gamma_\mu P^\mu)_{\alpha\beta} \quad (3.35)$$

$$\{Q_\alpha, Q_\beta\} = 0 \quad (3.36)$$

$$[P^\mu, Q_\alpha] = 0 \quad (3.37)$$

$$[P^\mu, \bar{Q}_\alpha] = 0 \quad (3.38)$$

The non-zero anti-commutation relation implies that one can produce spatiotemporal transformations by applying SUSY transformations. However, since all fields transform under space-time transformations, it follows that all fields will transform under SUSY as well.

With this mathematical machinery in place, one constructs supersymmetric theories, by imposing invariance under SUSY transformations on Lagrangians such as those of QED, QCD and so on. The fields $\psi(x_\mu)$ will no longer depend on the standard x_μ coordinates but also Grassmann number coordinates $\theta_\mu, \bar{\theta}_\mu$ acted on by the fermionic generators: $\psi(x^\mu, \theta^\alpha, \bar{\theta}^\alpha)$. Such a dramatic change calls for a new, cooler(!), name and our fields are now called **superfields!** Ultimately, one is able to introduce an extension of the standard model in which each fermion receives a χSF , each boson a $V SF$ and two chiral superfields for the Higgs bosons¹⁵. The partner fields will now form similar clusters to those we find in SM like leptons and quarks but labeled with the unimaginative addition of an s- prefix squarks, sleptons and so on - while counterparts to gauge bosons will be called gauginos. Much like the SM Lagrangian, the Minimal Supersymmetric SM (MSSM) Lagrangian will consist of four sectors: the gauge, lepton, Higgs and Yukawa. Most remarkably, however, the Higgs sector will now include fermionic fields as well.

The problem with MSSM is that the mass of the superpartners is the same as that of the SM fields. Since this is obviously inconsistent with what we (don't) observe in nature, we need to introduce a mechanism for breaking SUSY and giving the superpartners larger masses. To

¹⁵To avoid inflating this discussion any further we will skip the analysis of these terms. For more details one can consult (Signer 2009).

make this happen, we need to make sure that the vacuum is not invariant under SUSY:

$$\langle 0 | H | 0 \rangle_{SUSY} \neq 0 \quad (3.39)$$

This is done by appropriately modifying the potential V so that its minimum is positive: $V|_{min} = 0^{16}$. Since this does not work successfully for MSSM, one typically introduces an additional, hidden sector, which breaks SUSY in this sector and then connects to the visible sector of MSSM. The new, soft, terms included not only break SUSY but also result in new processes represented by Feynman diagrams such as the right diagram for the top quark (figure taken from Craig 2017, p. 20).




Figure 3.4: SUSY Cancelling Loops

The two processes cancel one another so that the quadratic dependence is eliminated (there are still logarithmic terms!). The same happens with the rest of the fields interacting with the Higgs and contributing to corrections of its mass. Clearly, naturalness can be a valuable guide for our expectations here: it will indicate the energy threshold at which new fields should be found if fine-tuning is to remain within acceptable bounds.

ASSESSMENT Despite the elegance of the above solutions and the fervor with which some, especially SUSY, were pursued, they suffer from the deadliest of syndromes a physical theory can suffer: lack of –or, arguably, even dis-confirming– empirical evidence. The novel physical predictions they make, such as the additional composite fields evoked in technicolor or the heavy superpartners in SUSY, have not been detected by the most recent runs at the LHC. While not explicitly ruled out, SUSY models come have under some severe pressure (Dine 2015, Giudice 2019). Even if SUSY is eventually vindicated through further observations, for it to count as a (good) solution to the hierarchy problem, the superpartners would need to be

¹⁶There are different ways to implement this: introducing so-called F -breaking or D -breaking terms or some combination of both types of terms.

found close “enough” to the 10^3 GeV energy scale in order to significantly reduce the fine-tuning required. While it is possible that SUSY symmetry breaking happens even at energies close to the Planck scale –so that failure to observe superpartners at lower energies does not amount to a falsification of SUSY–, its relevance to the hierarchy problem obviously diminishes with each failure to be confirmed. Discovering superpartners close to the Planck scale will do little for the hierarchy problem as the parameters will still need to be fine-tuned for the much lower scale of the SM Higgs. This indicates that naturalness by itself is not sufficient to boost our confidence in these alternatives to SM physics.

4. String Theory With the shortcoming of the above approaches noted, a radically different idea to find natural values for the Higgs mass has been gaining some momentum (Susskind 2008). Suppose that perhaps the problem has nothing to do with dynamical aspects of the theory, but rather concerns what is typical in a sample of possible worlds. What if there is a plurality of worlds and in each one the value of the Higgs mass is different? Given a sufficiently large set of worlds, it is certainly reasonable to expect that some of them will possess values compatible with our observations. The selective principle in this case is some form of anthropic reasoning. Since “the physical conditions necessary for our existence impose a selection effect on what we observe” (Smeenk and Ellis 2017, sect. 4.1):

Anthropic principle: “finely-tuned features of the universe [...] can be explained as necessary conditions for the existence of observer”

In accordance with this principle, the value of the Higgs boson mass will need to be compatible with existence of observers such as humans. Worlds, whose parameters are set to values that lead to life-hostile conditions, would preclude the existence of observers reporting on these values. And, incontestably, the fine-tuned value we have measured is favourable to (i.e. it is within some fortuitous range to render possible) the existence of life! Ultimately, then, there is nothing puzzling about the apparent unnatural value of the Higgs mass: it is just one of the life-friendly values in a vast space of possibilities. Would anyone be puzzled to find themselves living on solid land instead of the vastness of oceans?

Quite...naturally, the question emerges: what is the basis for lending credence to this extraordinarily bizarre scenario of a “plurality of worlds”? According to many physicists, physical grounds to support to this argument are found in the (happy?) marriage of string theory and inflationary cosmology. For years physicists working on string theory were frustrated by the

theory's failure to provide a unique vacuum solution. If string theory were to be the final theory, the holy grail of physics, the one theory to rule them all, how come it fails to specify in a unique manner the physical parameters, they wondered? As is frequently the case, however, the persistent bug was turned into a feature: the "indeterminacy" of the theory became a newfound source of strength. For advances in cosmology indicate the possibility of a mechanism for an interminable creation of parallel, bubble universes [Guth 2007]. Eternal inflation suggests that the process of inflation is ceaselessly operating throughout the whole universe constantly sprouting more universes in the ever-expanding sea of the multiverse. With this physical mechanism for the production of bubble universes, each equipped with its own initial conditions, the collection of string vacua of the string landscape could at long last escape the nebulous realm of mathematical potentiality to the physical world!

String Landscape + Inflationary Multiverse → Cosmic Lottery

There are two clear issues with this strategy. The first is of a technical nature. It has been emphasised that it is not sufficient to derive some possible solutions compatible with the observed values (if any exist at all). One also has to show these solutions to have a special status: being more typical than others. This resembles a transmutation of naturalness from claims about scale discrepancies in fundamental physics to selection problems within a multiverse setting (Williams 2019b). The conceptual bridge between these two seemingly distinct fields is the probabilistic style of reasoning (see figure 8.3).

Yet, as has been observed (e.g. Smeenk 2014, sec. 4, 5) forming meaningful probabilistic claims in this multiverse context -let alone obtaining accurate predictions- is a highly non-trivial task due to (at least) two conceptual-technical complications:

1. One has to first of all define a sample space and a **measure** to even render the use of probability distributions meaningful. It is far from clear that a well-defined space exists or that a unique measure can be picked out in a non-arbitrary manner.
2. Even if a probability distribution can be defined, it remains far from obvious whether it can be accorded some physical significance. For example, a standard way to connect these distributions with physical processes in statistical mechanics is to invoke ergodicity¹⁷, but is far from clear how something similar can be achieved in a cosmological setting.

¹⁷Ergodicity, the assumption that all microstates are equiprobable over sufficient long time scales, essentially allows us to make a correspondence between the phase space of the system and its physical instantiation by mapping the volume of regions of phase space to the time a system will remain at a given state.




Figure 3.5: The evolving conception of naturalness

The second issue is conceptual and perhaps even more unsettling. By shifting our attention from naturalness as a problem about scales to naturalness as a problem about the typicality of certain values for physical parameters, we might be missing on the crucial physical insight that this principle was (and is?) meant to encapsulate. In this manner the demand for naturalness acquires a more heuristic character perhaps making criticisms, such as those voiced by Grinbaum 2012 (p. 627), particularly apt:

As in the general case of probabilistic reasoning in a situation of uncertainty..., the fine tuning argument is the last resort of the mind when no rational guidance to future results can be provided. [...] However, the principles of nature, both known and unknown, are unique and unstatistical. Therefore, there is no firm epistemological ground to believe that fine-tuning actually leads to a true theory.

Of course, this is no reason to reject this formulation of naturalness, but one must lose sight of the conceptual shift that has taken place here. Naturalness is now a standard on which one can compare theories – perhaps a standard that can be added to a catalogue of theoretical virtues like simplicity, fruitfulness etc. As a consequence, the only acceptable “justification” will come from its capacity to lead to the adoption of more empirically successful theories.

3.3.1 Prediction or Reconstruction?

The way we have presented it, the significance of the hierarchy problem heavily depends on whether we accept naturalness in some form or another. Unsurprisingly, then, we need to turn to an assessment of the latter’s status. Consistent with our discussion of a more “axiological” form of realism in the previous chapter we want to examine whether problems based on naturalness are worth pursuing. Lacking any “raw material” of the experimental kind, history must be our guide here. What we need to do is assess to what extent naturalness can be seen as a guiding principle (actively) contributing to the development of particle physics or whether it is more

akin to a post hoc reconstructive principle. Following Sahuquillo 2019 we may distinguish between two applications of a principle like naturalness:

- **Reconstructions:** counterfactual statements of the form “if principle X was employed in context-problem C, it would have led to predictions P”
- **Predictions**¹⁸: factual statements of the form “the principle X was applied to problem C and led to prediction P”

At first glance, this might strike someone as a rather pedantic distinction. However, the history of science is fraught with examples of principles that were at one point taken to encapsulate deep insights about nature, but later found to be lacking. One can think of the generalisation of the principle of relativity, which was initially thought to be *the* physical basis of general relativity (whence its name). From a pretty early point it was seen that Einstein’s goal of treating accelerated on equal footing with inertial systems did not exhaust the content of GR. The general principle of relativity would need to be supplemented with further assumptions to become substantive. (see Norton 1993, Pooley 2015). However, there also exist principles - physical insights that can be used to reconstruct the rationale behind particular theoretical developments. Sometimes, however, such reconstructions rely on some form of equivocation or a more idiosyncratic reading of history. One can think of Rovelli 2007 (chapter 1) and his reconstruction of the development of spacetime theories as a progression towards the elimination of superfluous structure – with the obvious endpoint being the elimination of background spacetime itself, a fact that can be understood as an advantage of LQG over alternatives like string theory. Such principles can only be said to have vaguely or heuristically guided theory development, perhaps along the way concealing the true facts that led to or rendered those breakthroughs possible¹⁹. Since even predictive principles occasionally fail in shedding light into the unknown regions of future physics, reconstructive principles are even less trustworthy in this respect.

For our project it will be sufficient to examine specific historical episodes where naturalness violations led to the formation of correct expectations for future theories. If naturalness can be

¹⁸Sahuquillo uses the term “success” but for our purposes “prediction” is more convenient to stress the contrast between the active or passive role a (rationalistic) principle like naturalness can play. Prediction, as opposed to mere reconstruction, involves some epistemic risk.

¹⁹This is, of course, not meant to diminish the role of such reconstructions. One can definitely gain a deeper insight and better appreciation of a theory’s novel features through a careful counterfactual analysis of history. This is the case with the counter-factual story Stachel 2007 gives about Newstein and the role of differential geometry tools in the making the leap towards the more accurate theory of gravity.

explained away or was simply missing in these alleged positive instances, its status will have to be reduced to that of a rational reconstruction principle.

1. As we saw, Dirac deemed the huge discrepancy between the proton mass and Planck scale a puzzling fact begetting some explanation. Perhaps his way of addressing the problem, with the introduction of new cosmological ideas, was misguided, but accepting the puzzlement as a genuine problem arguably did lead to the right explanation... in particle physics! Indeed, dimensional transmutation in QCD, a procedure with which a dimensionless parameter like the ratio of two scales is turned into a dimensionful parameter, shows why the ratio between the two scales is so small. Solving the relevant RG equations (connecting the QCD scale and the Planck scale) one can show how this follows from confinement:

$$\Lambda_{QCD} = M_{Pl} e^{-\frac{2\pi}{7\alpha_3(M_{Pl})}} \quad (3.40)$$

where M_{Pl} the mass of scale associated with the Planck scale and Λ_{QCD} the scale at which QCD breaks down. Since we have an exponential decay, the massive discrepancy between the scales is to be expected. Note that the naturalness at play here is Dirac rather technical naturalness; we did not invoke any particular “protector” symmetries.

2. Similarly, Craig 2017 treats the example of scale discrepancy (of order 10^{-5}) between leptonic masses (such as the electron and muon) and the top quark as another example of a successful application of naturalness. Briefly, the answer for the discrepancy derives from the fact that flavour is a technically natural symmetry. When the Yukawa couplings in the Lagrangian are taken to zero the theory’s symmetry is augmented to a $U(3)^5$ group. When they are turned on corrections will be proportional to these very small couplings of order 10^{-5} . While the small ratio of the two scales might still seem a problem worthy of an explanation, technical naturalness at the very least guarantees that since the corrections are proportional to the parameter, no fine-tuning problem will arise. That’s because the corrections will be proportional to the parameter and thus the latter will drag down the former when set to pretty low values.
3. Another application of naturalness might be seen to be the discovery of the ρ -meson. Postulating the existence of the meson can be seen as a “response” to the unnatural dis-

crepancy in the masses of charged and neutral pions:

$$m_{\pi^+}^2 - m_{\pi^0}^2 \sim (35.5 \text{ MeV})^2 \quad (3.41)$$

The key fact is that electrically charged pions can interact with the EM field and consequently their mass will receive corrections from photon loops: $\delta m = -\frac{3e^2}{16\pi^2}\Lambda^2$. This result will need to match the experimentally measured value above and the natural way to do so is to postulate some new propagating particle close to the energy scale difference, 850 GeV. The ρ -meson, with a mass around 770 GeV vindicated this expectation.

4. The strongest case for the utility of naturalness is probably the discovery of the charm quark in the 1970s. The puzzling scale discrepancy in this case was the difference in the mass of neutral and charged kaons:

$$\frac{m_{K^+} - m_{K^0}}{m_{K^+}} \sim \frac{G_F^2 f_K^2}{6\pi^2} \sin^2 \theta_c \Lambda^2 \quad (3.42)$$

with G_F the Fermi constant, $f_K = 114 \text{ MeV}$ the kaon decay constant and $\sin \theta_c = 0.22$ the Cabibbo angle. This calculation was performed at one loop using the $V - A$ (vector minus axial) theory available at that time, which was one of the first attempts to describe the weak interaction in the QFT framework. The scale Λ appears explicitly here because this theory is non-renormalisable - it is an effective theory for the more fundamental renormalisable electroweak theory that we currently possess.

What was the problem, then? The experimental value for the ratio is known (today) to be around 7×10^{-15} . The above calculation is part (actually, the first order) of a correction expansion that includes further terms that are not shown above. Therefore, since some contributions are omitted, the theoretical value will need to be somewhat smaller than the measured one. This gives an estimate $\Lambda < 2 \text{ GeV}$. There are two alternatives for the $V - A$ theory:

- It remains valid up to higher energies. Then, assuming that these energies are some orders of magnitude above 2 GeV, one would need to fine-tune the coupling constant to make the theory agree with the experimental result.
- It gives way to some new theory at an energy scale low enough to eliminate the necessity to fine-tune the parameter.

Glashow, Iliopoulos, and Maiani 1970 took the unnatural restriction on the cut-off of the theory to be problematic and suggested a mechanism, the GIM mechanism, that postulated a new quantum number, charm, along with a new quark, the charm quark. This new field also contributed to the above corrections by cancelling internal up quark lines. Comparisons with the measured value for the kaon mass difference led to estimates of the charm quark mass at around $m_C \approx 1.5 \text{ GeV}$ ²⁰.

All these examples provide nice illustrations of how (for the most part) technical naturalness is supposed to work in action. Unfortunately, with the exception of the charm quark discovery, most of them do not constitute real historical episodes, but later reconstructions of these episodes using naturalness-style reasoning. No one invoked naturalness to predict the ρ -meson discovery nor the positron. The way we reasoned from the unnaturalness of certain parameters to the postulation of some new structures was not the actual road taken. To make a stronger case, however, we will soon turn to the issue of whether these post hoc rationalisations should be seen as vindications of the principles they meant to exemplify.

Bad news have recently come from the allegedly well-established case of the charm quark, though. Recent work by Sahuquillo 2019 has questioned the standard appraisal of above the story. His main contention is that, since naturalness was not yet properly formulated before the late 1970s, the idea of attributing the discovery of the charm quark to its violation is in the final analysis anachronistic (p. 59):

we have shown that there is a certain rationale, as stated by Gaillard and Lee, that could be effectively addressed as naturalness assumptions. However, it is important to notice that it is not the only argument used in the paper in order to perform the computations. Indeed, the main motivation of the paper is to explore the Kaon decays within the Weinberg-Salam model with four quarks, thus introducing the GIM mechanism within a renormalizable model of (electro)weak interactions.

Surveying the rest of the papers engaging with this problem, Sahuquillo demonstrates that, even though not completely extraneous or inapplicable, naturalness considerations were neither the primary driving force behind the postulation of the charm quark nor were they applied explicitly constraints on model construction. Summarising the above results:

²⁰Wondering how exactly? The charm mass quark is assumed to be some factor modifying the coupling constants on the RHS of equation (8.42). By comparing this theoretical calculations to the actually measured value one can fix the kaon mass so that the measured and theoretical value agree.

Table 3.1: Naturalness Record

Success	Failure
<ul style="list-style-type: none"> • Electron scale* [CURE: positron] 	<ul style="list-style-type: none"> • Higgs Boson mass [conjecture: technicolor (refuted?), SUSY, extra dimensions]
<ul style="list-style-type: none"> • Charged vs Neutral Pions scale discrepancy* [CURE: ρ-meson] 	<ul style="list-style-type: none"> • Cosmological Constant [conjecture: string landscape-multiverse?, SUSY?]
<ul style="list-style-type: none"> • Neutral Kaons mass discrepancy [CURE: charm quark] 	<ul style="list-style-type: none"> • Strong CP [conjecture: axions]

Verdict There are as many potential successes of naturalness as there are failures. The problem with the alleged successes is that a) all but one case constitute post hoc rationalisation of the discoveries and b) even the single case where naturalness was invoked to justify the expectation for a new particle its contribution is of dubious status. Even if this true, however, why on Earth can't reconstruction principles of a reconstructive kin guide our expectations for future developments? Our conceptual understanding, the dissenter continues, was not mature enough to appreciate the principle at work back in the 1970s, but as soon as naturalness was grasped, its fundamental role as a precondition for particle physics could be projected (better: mapped) onto key episodes of theoretical breakthrough. That is a fair point to raise, but still fails to be fully convincing for the simple Humean predisposition of the (human) mind to spread itself on the world. Specifically, The post hoc character of the principle is relevant because it is not unlikely that the regularity we observe is a regularity "imposed" by us on the development of SM physics as some sort of aesthetic predisposition rather than a true regularity derived from experience. Similar to Whiggish historians, we turn to the history of science to vindicate our agenda tailoring the facts in a manner that will suit a pleasing narrative. We all know that this project usually does not fare well...

Whiggish historical narratives aside, a vivid analogy to appreciate the subtle difference between predictions and reconstructions comes from economics! Despite the popularity eco-

conomic studies have been enjoying in modern times, the field is singular in that clashes over its basic assumptions have sent pretty erosive shockwaves to its very foundational core. Prominent economists have oftentimes disputed the scientific status of economics questioning the very foundations of the currently prevalent economic paradigm. What is the strongest piece of evidence in favour of these dissenting voices? The case is clear and pretty appealing: economists have failed to predict crises as important and devastating as the 2008 global financial crises. In fact, things are far worse than that: the very assumptions behind the way financial markets were taken to operate appeared to rule out the possibility of the massive breakdown the world experience about a decade ago. Clearly, then, some things must change...

But, what exactly, is the connection with our preceding discussion of naturalness and reconstruction? Well, think of the following (overly) simplistic example. Let us assume that Milton and his daughter Ayn are playing some card game. At a crucial juncture in the course of the game, Milton plays a superb card combination and effectively cements his position as the winner. Now, this makes perfect sense from the neoclassical theory: Milton, as a rational agent, sought to maximise his utility function and played the strongest card combination possible to do so. Now, imagine that Milton refrained from playing that combination, perhaps going as far as playing a pretty bad one, and went on to lose the game. Did Milton act in an irrational manner in this case? Well, this conclusion is most definitely not forced on us²¹. Instead, what an economist studying this toy example could say is that Milton's ranking of preferences is different than what was assumed: Milton actually did not want to win the game, but merely present enough of a challenge for his daughter to enjoy it. The ultimate goal, however, has always been to let her win. If this does not already strike one as problematic, consider the following twist for the first case: Milton did not have any true interest in winning the game but he wanted to teach his overconfident daughter a lesson in humility or, to use a less ascetic motive, simply hone her gaming skills and therefore played his best to make his case.

Now, what exactly is the problem with this? Well, if you think about it, with enough imagination and resourcefulness one can come up with a principle to rationalise any sort of behaviour. One can thus render the theory compatible with any observed behaviour; any failures of prediction are attributable to a poorly conceived set of prior preferences. The framework can be of reconstructive use but ultimately it is completely unsatisfactory as a tool for pre-

²¹And it shouldn't! Imagine that Milton was playing cards with his pal Friedrich the previous day and played the exact same combination to brutally shatter his hopes of turning around the game. Are we to assume that Milton's rationality was compromised within the course of a day?

dicting, let alone truly understanding, human behaviour. I claim that something analogous is true of reconstructive vs predictive principles: in the absence of the epistemic risk involved in guiding actual theoretical developments any carefully chosen principle could do the trick. We simply cannot be certain of the assumptions that played an indispensable role in those episodes examining them purely in (philosophical) *vitro*. The problem is that, paraphrasing and re-engineering Kant's point, without experience principles are blindfolded: in vain do we grope in the (empirical) dark for their intentional object²².

All things considered, naturalness may not be vindicated through actual scientific practice. At best, it constitutes a reconstructive principle that can be superimposed onto the actual historical development of QFT and particle physics. We need to turn to alternative sources of motivation. This is exactly the task that we will take up in the following section after examining some arguments against the coherence or usefulness of the notion.

3.4 Reaction 2: Denial

We now turn to the skeptical reactions towards the hierarchy problem and the status of naturalness as a guiding principle for (future) physics. We will present the challenges in three waves. First, some technical objections about the choice of renormalisation and regularisation schemes and their corresponding physical significance. Then, some skeptical remarks about the mathematical well-definiteness and persuasiveness of fine-tuning arguments. Finally, we attempt to unearth the physical insight - motivation lurking beneath naturalness.

3.4.1 Technical Objections

Physicists and philosophers alike have disputed the cogency of the hierarchy problem as standardly presented by insisting that it relies on some confusion or misunderstanding of the technicalities or physical significance of the process of renormalisation. People arguing along these lines typically contend that the problem disappears when an appropriate regularisation scheme is chosen (implicitly assuming of course that physical results need to be scheme independent) or goes up in smoke when a renormalisation condition is eventually imposed.

²²Lest the analogy with economics be seen as beneath a pure-blooded science like physics, let us refer to the principle of equivalence as a concrete example of such a retro-active principle. Indeed, it eventually became clear that PE did not uniquely specify GR as the correct theory of gravity. Various (weak, Einstein, strong) formulations of the principle have thus been given with the goal of narrowing down the space of alternatives.

Bug #1: Regularisation The possibility of attributing the problem to a poor regulator choice has been suggested by various sources. In their brief introductory textbook on QFT, Alvarez-Gaumé and Vázquez-Mozo 2011, for example, already note (p. 236):

The fact that DR eliminates quadratic divergences might seem surprising in the light of the previous discussion of the hierarchy problem. Indeed, as DR regularizes the quadratic divergences to zero it seems that the whole hierarchy problem results from using a clumsy regulator, and that by using DR we could shield the Higgs mass from the scale of new physics.

However, they go on to immediately note that:

This is not the case, but for interesting reasons. In spite of DR the Higgs mass is still sensitive to high energy scales.

Manohar 2018 echoes the first sentiment and strikes a fairly dismissive note against the problem (my additions in brackets):

You will have heard endless times that Fig. 5.2 [the loop correction to Higgs mass diagram] gives a correction: $\delta m_H^2 \propto \Lambda^2$, to the Higgs mass that depends quadratically on the cutoff. This is supposed to lead to a naturalness problem for the SM, because the Higgs is so much lighter than Λ , which is taken to be at the GUT scale or Planck Scale. [...]

The above argument for the naturalness problem is completely bogus. The regulator used for the SM is dimensional regularization, which respects gauge invariance. The actual value of the integral is eqn (5.11) [which involves no cut-off regulator]. Adding the renormalization counterterm cancels the $\frac{1}{\epsilon}$ piece, resulting in a correction to the Higgs mass $\delta m_H^2 = -12\lambda m_H^2 \log \frac{m_H}{\mu^2} + 1$ which is proportional to the Higgs mass.

Doubling down on this dismissive attitude he devalues the hierarchy problem to a mere fixation with coincidences by comparing it to the following example:

Finally, let me comment on another fine-tuning problem that many of you are excited about. There will be a total solar eclipse on Aug 21, 2017, shortly after the Les Houches school ends. The angular diameter of the Sun and Moon as seen from Earth are almost identical—the Moon will cover the Sun, leaving only the solar corona visible (see Fig. A.1). The angular diameters of the Sun and Moon are both experimentally measured (unlike in

the Higgs problem where the Higgs mass parameter m at the high scale M_G is not measured) and the difference of angular diameters is much smaller than either. Do you want to spend your life solving such problems?

Echoes of this last point can be heard in Rosaler and Harlander 2019 who compare the problem to a poor choice of coordinates for the measuring of distances (p. 131):

With choice of origin at the center of the Milky Way (O2), there is an extremely delicate cancellation - to roughly one part in 1018 - between vectors indicating the locations of Aachen Dom (Cathedral) and Aachen Rathaus (Town Hall). However, this cancellation can be dramatically reduced by moving the origin close to, say, the midpoint between the two (O1). Likewise, the cancellation between the bare Higgs mass and its quantum corrections is very delicate for large Λ (say, the Planck scale), and much smaller for Λ on the order of the Higgs pole mass. Without fundamental parameters, the choice of the unphysical reference scale Λ is purely conventional and akin to a choice of origin; the delicate cancellation between bare Higgs mass and quantum corrections is then an eliminable, unphysical artifact of convention.

In both (dismissive) discussions of the naturalness, the problem, of delicate calculations with the higher energy regime to produce accurate results at the lower energy scale, is attributed to an inconvenient parametrisation scheme. Set out to measure the width of your favourite QFT book using your distance from the nearest physics department as your adopted length scale and you are bound to need all sorts of weird fine-tuning adjustments in last digits to produce an accurate result.

It is worth examining the more technical aspects of these claims in some detail. Recall that the quadratic divergence problem is essentially the quadratic dependence of loop corrections to the Higgs boson on the cut-off scale:

$$\delta m_H^2 = g \frac{\Lambda^2}{16\pi^2} + \dots \quad (3.43)$$

Obviously, the Λ dependence is only significant in a cut-off regulator scheme and will not come up if one applies dimensional regularisation in its stead. In this case, the corrections are rather of the form:

$$\delta m_H^2 = g \frac{m_H^2}{16\pi^2} \left(\frac{1}{\epsilon} + \frac{1}{2} + \ln \frac{\mu^2}{M^2} + \dots \right) \quad (3.44)$$

with no dependence on anything like the hard cut-off Λ .

Proponents of naturalness typically insist that this does not help assuage the worry because the problem is not the dependence on the cut-off per se but rather the apparent influence of higher-energies to lower-energies. The problem is that when we adopt a mass-independent scheme like dimensional regularisation, we will eventually need to match the low and high energy theories at $\mu = M$, i.e. the scale above which processes are suppressed. When this is done, the mass term in the Lagrangian will be shifted to:

$$m_H^2 = m_0^2 + \frac{g}{4\pi^2} M^2 \quad (3.45)$$

Since the processes that were supposed to be suppressed in the EFT come back to haunt us by contributing to the value of the effective parameter, we have been unable to eliminate the sensitive dependence of low energy physics on high energy physics. Heavy fields do not decouple. Note that m_0 will again need to be chosen in a way that will delicately cancel the high energy contribution²³.

Bug #2: Renormalisation A more recent charge against the significance of the hierarchy problem has sprung from its (alleged) disappearance when an appropriate renormalisation scheme is selected. Both Bain 2019 and Harlander and Rosaler 2019 have argued that the adoption of the on-shell mass renormalisation scheme leads to a perfectly finite result for the mass of the Higgs boson. Indeed, as one can see in Schwartz's (2014, p. 408-9) toy theory of fermion-scalar coupling, after applying the (standard) on-shell conditions:

$$\delta_m = \frac{1}{m_P^2} \Sigma_2(m_P^2) \quad \delta_\phi = - \left. \frac{d\Sigma_2(p^2)}{dp^2} \right|_{p^2=m_P^2} \quad (3.46)$$

with m_P the pole mass, one evaluates the loop correction at

$$\Sigma(p^2) = \frac{\lambda^2}{4\pi^2} \left[\frac{(p^2 - m_P^2)^2}{20M^2} + \mathcal{O}\left(\frac{m_P^6}{M^4}\right) \right] \quad (3.47)$$

²³A note on how fermion fields manage to evade this kind of trouble. Let us reverse the assumptions of the toy example and take the fermion mass m_f to be much smaller than the mass of the scalar field m_P . As a result, corrections to the fermion mass will only be proportional to the fermion mass itself and will not depend on the heavier scalar mass. Clearly, technical naturalness is satisfied in this case. More importantly though, the example serves as an illustration of the connection between technical naturalness and what can be seen as the most disconcerting feature of the hierarchy problem: failure of decoupling.

where it is clear that as the fermion mass M becomes higher its contribution becomes lower (scaling as $\sim \frac{1}{M^n}$).

Assessment: There are two problems with a dismissal of the problem along these lines, however. First of all, the above points do not speak against the third formulation of the problem which, recall, concerns the behaviour of the mass parameter under RG transformations. The choice of scheme, even if it resolves issues of fine-tuning, does not really touch upon the third issue: that the initial values for the parameters need to be finely adjusted to derive the correct macro-physics. One possibility here is to rethink the exact relation between the RG trajectories and EFTs following an approach suggested by Rosaler and Harlander 2019. According to them, we should drop the assumption that a theory is represented by a point in parameter space²⁴. and associate it with a whole RG trajectory, instead. This new take on the RG space implies that RG trajectories will no longer represent distinct theories across different energy scales, but will rather comprise something akin to equivalence classes: choosing a particular scale will more or less be like choosing a parametrisation carrying no physical significance. Undoubtedly, this proposal, as Rosaler & Harlander themselves admit, is still tentative; it is not yet clear that their suggestion can be made to work at a technical level. For example, a theory requires a Hilbert space on which observables are meant to act. But what is the Hilbert space for a theory whose observables run wild across different scales? Presumably one would have to indicate how the “put-in-by-hand” procedure of the continuum group is to be included in definition of an EFT. That is not all, though. Another, particularly significant, source of concern is that the abandonment of the received understanding of the connection between RG and EFTs (which we explored in some detail in chapter 4) seems to push us back to a previous stage of our understanding of the process of renormalisation. In effect, we relinquish our reading lower-order processes as effective descriptions of whatever underlying fundamental theory in favour of the admittedly less insightful “cancellation of infinities” story. The strong analogical ties with statistical mechanics and the coarse-graining of smaller-distance degrees of freedom is severed.

The significance of the above notwithstanding, there is a (more philosophical) price to be paid should the problem be dismissed as a mere technicality. When using an on-shell renormalisation scheme we fix the renormalised mass to the value of pole mass, which is found exper-

²⁴Recall that each point corresponds to a set of specified parameters at different energy scales.

imentally. This amounts to the introduction of a top-down constraint as a means to accurately fixing physics at a more fundamental level. Is this compatible with the reductionist program that has dominated and guided particle physics throughout the 20th century (e.g. Weinberg 1987)? Prima facie this looks like a bitter pill to swallow. The suspected readers will probably have guessed how our philosophical discussions in the previous two chapter will bear on the stance we take here. Nevertheless, to keep them in suspense, we will postpone this discussion for the last part of this chapter. This will give us time to complete our cartography of the whole problematic before drawing our final conclusions.

3.4.2 Disputing Fine-Tuning

The most popular way of presenting the hierarchy problem is as an instance of fine-tuning in need of explanation. The bare mass or matching condition or initial conditions need to be very very very precisely calibrated to lead to descriptions matching experience. Surely, we cannot find that much epistemic luck reasonable! Undeniably, for this argument to have real bite it is crucial to substantiate (or lest this be too strong a requirement to clarify, make more intuitive) the kind of probability claim made here. What is the meaning of a statement like “it is unlikely that the value of parameter λ_ν be in the interval $I = [a, b]$ ”? For a well-defined answer one needs to a) specify a procedure for making a choice (the “lottery”) and b) define the notion of likelihood employed (the “chance”).

Consider the following scenarios:

1. We cast a die once and obtain a “6”. We cast it a second time and obtain a “6” again. We repeat the same process a hundred times and obtain “6” in each repetition.
2. We put a large ball on the top of the tiny summit of a curved surface and wait until the ball starts rolling. The ball stays there for hours.
3. We balance a pencil on its tip and it stays there for a few milliseconds. We try again but this time it remains standing for a whole hour.

Intuitively, there is something deeply unsettling about all these scenarios: the specified configurations and/or their “lifetimes” are highly atypical. Normally, we would expect the die to land on a different number at least once over a hundred repetitions, the ball to start rolling down the surface quickly after putting it on the top and the pencil to almost immediately land on one

of its sides as soon as we cease holding it. It has long been a problem for statistical mechanics to connect the dynamical evolution of a system with the number of possible configurations of its constituents. The central idea is that the typical or most likely to obtain macroscopic states for a system are those that correspond to the largest number of configurations at the microscopic level. These typical states are identified with the equilibrium states we come across in thermodynamics and the study of macroscopic phenomena. It was Boltzmann's great insight to use this identification to explain the tendency thermodynamic systems appear to have towards equilibrium. Other things being equal, a system's micro-configurations will ultimately evolve to a configuration corresponding to a typical, equilibrium macro-state. To return to our examples, it is obvious that obtaining a hundred sixes in a row is atypical for a fair die; such an unlikely string of results! It does not call for a wary temperament to suspect that the die is rigged.

Can we follow a similar rationale in the context of QFT in order to put the aversion to fine-tuning on a precise ground? If we think of the parameter space as analogous to an event space with various regions corresponding to different choices for physical parameters -and by extension different physics-, perhaps we could accord typicality (naturalness) a quantitative meaning by measuring the volume of regions compatible with the correct (empirically verified) physics. If the regions of parameters that lead to the world as described by the standard model are very small²⁵, we will be able to claim that the specified set of parameters is highly untypical and, as a result, the theory unnatural. Wallace 2019 seems to be thinking in terms of this analogy when he compares the notion of naturalness in particle physics with that of a typical distribution in statistical mechanics. We will examine his argument in more detail in the following section. For the moment we will only discuss the technical aspects of defining an appropriate distribution.

Measures Much like in statistical mechanics, the main problem is singling out a measure over the sample space. Unfortunately, unlike statistical mechanics, where choosing the Liouville is justified by the fact it remains invariant under the dynamics, there is no obvious choice to be made in this context. This is a huge problem because the choice of a particular measure will determine the degree of fine-tuning required. To understand this, consider the example

²⁵Because, if the bare Higgs mass or the cancellations between renormalized parameters and their corrections need to have very specific values, then there will only be a small set of values that will be compatible, especially with accuracy close to 10^{-18} .

presented by Wells 2019: let X_1, X_2 be two random variables that follow a uniform distribution over $[0, 1]$ and define the new random variable $Y = X_1 - X_2$. Recall that to find the probability distribution for Y we need to find the convolution of the probability densities of X_1, X_2 :

$$f_Y(s) = \int_{\mathcal{R}} f_{X_1}(s - y)f_{X_2}(y) dy \quad (3.48)$$

Substituting for the densities of x_1 and x_2 :

$$f_{X_i}(s) = \begin{cases} 1 & s \in [0, 1] \\ 0 & otherwise \end{cases} \quad (3.49)$$

we obtain the density for their difference and integrating it up to value s we get the probability distribution $P(Y < s) = F_Y(s)$:

$$F_Y(s) = \begin{cases} 1 + s & s \in [-1, 0] \\ 1 - s & s \in [0, 1] \\ 0 & otherwise \end{cases} \quad (3.50)$$

from which it becomes clear that the most likely values for the difference between the two values is close to 0. Contrary to our intuition, fine-tuned values for Y are highly typical outcomes. This drastically changes, if we adopt a different measure for X_1, X_2 , however. For example, assuming that the two variables are independent and normally distributed over $(0, 1)$, $X_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$, their difference will also follow a normal distribution with $\mu = \mu_1 - \mu_2$ and $\sigma = \sigma_1^2 + \sigma_2^2$:

$$Y \sim \mathcal{N}(0, 2) \quad (3.51)$$

and small values for the difference will be assigned very low probabilities and therefore will constitute highly atypical results.

Hossenfelder 2018 has forcefully pressed this point questioning the implicit choice of a uniform measure over the parameter space. At a superficial level, the choice of a uniform measure should be a no-brainer. After all, what better choice do we have than to treat all possibilities on equal footing? The latent assumption here is an appeal to some form of indifference principle:

Principle of Indifference: In the absence of relevant evidence in favour or against some possible outcomes (events) under examination, an agent should distribute

their degrees of belief (or credences) evenly among them.

In other words, when lacking evidence to weigh events one way or another the PoI prescribes that we assign them equal probabilities. Since none of us had had a chance to consult with the Almighty during the early days of creation, we are left with no other choice but to treat the problem of assigning probabilities to these physical parameters as an instance of reasoning in ignorance. But how can something come out of nothing? Wouldn't it be more prudent to suspend judgment in light of the lack of evidence to sway our thought towards one side or the other? In fact, critics have long ago stressed this point. The principle of indifference can be shown to be highly dependent on the exact formulation of the problem²⁶. This means that different ways of posing the question, i.e. setting up the sample space, will lead to widely varying estimates for what should be considered a typical vs an atypical outcome. Given this indeterminacy, one must be vigilant in their application of PoI. Otherwise, they are bound to be in the unpleasant position of Hossenfelder's (Ibid, app. B) pounded pro-naturalness interlocutor:

To better see why this criterion is circular, think of a probability distribution on the interval from 0 to 1 that is peaked around some value with a width of, say, 10^{-10} . "There," you exclaim, "you have introduced a small number! That's fine-tuned!" Not so fast. It's fine-tuned according to a uniform probability distribution. But I'm not using a uniform distribution; I'm using a sharply peaked one. And if you use this distribution, then it is very likely that two randomly selected numbers are at a distance of 10^{-10} . "But", you say, "that's a circular argument." Right, but that was my point, not yours. The sharply peaked probability distribution justifies itself as much or as little as the uniform distribution does. So which one is better?

Simply wind: sow the wind, reap the whirlwind. Intuitive choices sometimes lead to very puzzling results. It is worth noting that Anderson and Castano 1995), among the first to suggest measures for fine-tuning, already underlined that "this 'theoretical license' at one's discretion when making this choice necessarily introduces an element of arbitrariness to the construction" since "in the absence of a theoretical reason compelling us to choose a certain value, we can consider some sensible distribution of the parameter to study what are the natural predictions of the model" (p. 302). There exist numerous suggestions for measures of fine-tuning in the literature²⁷

²⁶A fact specifically underlined by the so-called Bertrand's paradox. Check chapter 3 in Myrvold 2021 for a more detailed presentation of this argument and its consequences for the PoI.

²⁷Perhaps one of the most "intuitive" measures is the one proposed by Barbieri and Giudice 1988, which

Bayesianism Since the above theme failed to inspire, it might be time to turn to a new tune. It has been suggested (e.g. Fichet 2012) that the problematic around naturalness can be put on firmer ground if viewed through the lens of Bayesian confirmation theory. Indeed, Bayesian probability theory seems to be the perfect framework in which we can address questions of the form: “how likely is that the standard model is true at low energies given a higher energy theory”? Usually this can be represented as $\mathcal{P}(SM|BSM)$. Before shifting to a formulation of the hierarchy problem in this frame, we will provide a very brief sketch of the Bayesian approach:

Bayesians apply a central result of the theory of probability, Bayes’ theorem, which connects the probability of an event A conditionalising on an event B :

$$P(A|B) = P(B|A) \frac{P(A)}{P(B)} \quad (3.53)$$

to credences, i.e. probabilities as representing the degrees of belief of a rational agent. This allows us to make estimates about the confidence an agent possessing information I can have on a hypothesis H in light of some (empirical) data D :

$$P(H|I, D) = P(H, I) \frac{P(D|H, I)}{\sum_H P(D|H, I)P(H, I)} \quad (3.54)$$

with $P(H, I)$ called the prior probability, that is, the probability representing the agent’s degree of belief in hypothesis H before updating their knowledge with the new data D . Unsurprisingly, $P(H|I, D)$ is the posterior probability. The ratio in the RHS encodes the “epistemic impact” of the data to the agent’s credence.

Using this device one can not only talk about the probability a certain hypothesis is true under a set of data, but also compare two models M_1, M_2 on this very basis. One simply writes the expressions for (7.51) for the two models, eliminates the common term $P(D)$ and obtains a comparative ratio:

$$\frac{P(M_1|D)}{P(M_2|D)} = \frac{P(D|M_1) P(M_1)}{P(D|M_2) P(M_2)} \quad (3.55)$$

compares changes of an input parameter at higher energies \mathcal{X}_i to some fixed mass μ scale like that of the Higgs boson:

$$FT[\mathcal{X}_i] = \left| \frac{\partial \log \mu^2}{\partial \log \mathcal{X}_i} \right| \quad (3.52)$$

The rate of change of the low energy scale parameter to changes at higher energy scales is an indicator of the sensitive dependence of low energy to higher energy physics. Anderson and Castano 1995 revised this measure by dividing it with an average fine-tuning in order to turn into a tool of theory comparison. Williams 2019b nicely tracks how these modifications led to a novel form of naturalness disentangling naturalness from sensitivity.

The first ratio in the RHS codifies the relative degree of belief between the two models. When it is > 1 it will make $P(M_1|D)$ higher and conversely, when it is < 1 it will do so for $P(M_2|D)$. One can thus compare models with different parameters

Setting the priors to a uniform probability distribution appears to be a reasonable move as it truly represents the maximum amount of knowledge we can have when lacking the evidence to impose some ranking. There are two problems with this. First, while appealing, the choice of a uniform distribution for the priors still falls prey to the same problem we saw above: “an absence of any known reason to judge one alternative more probable than another (which would warrant at best a suspension of judgment), and a *positive judgment* of equiprobability’ (Myrvold 2021, 3.1.3, my emphasis). Even if someone is skeptical of this distinction between suspension of judgment and positive judgment, it is still worth asking themselves: isn’t the fact that the choice of a particular distribution renders a highly successful theory like SM unlikely all the more reason to reject it? How can we assign SUSY a higher degree of belief than the SM when we simply lack any experimental evidence in its favour? Quite a reversal of fortune!

Perhaps this remark only speaks against an absolute notion of naturalness (as in: model M is likely or unlikely), but leaves its relative - comparative aspect [as in: model M is more likely than model N] unscathed. Even if this is true, and setting any technical complications aside, the conceptual issue we noted above has now become even more pressing. We have shifted our attention from naturalness as a physical(ly motivated) principle to naturalness as a criterion - principle for theory selection. Williams 2019b (see figure 3.5 again) explains the current trend shift in discussions of the multiverse, the string landscape and typicality as evolving branches of this (quite dissimilar) problematic about naturalness in mid-late 20th century.

In the end, one cannot help but raise the sceptical doubt: if naturalness as an extra-empirical principle for theory selection leads to results incompatible with standard empirical methodology *and* the conceptual ground of fine-tuning arguments is (at best) shaky, how seriously should we take this reformulation of naturalness along confirmation theory lines? If anything, the style of reasoning that leaves the SM as a “less likely” theory in spite of its tremendous empirical success in favour of some at best tentative theoretical speculation should be read as a warning sign. Perhaps it is high time our directives were reassessed.

3.4.3 Seeking Motivation

Let us try to “catch a breath” over the following question: what exactly is the notion of naturalness meant to encode? Put differently, what is the more profound underlying physical insight that finds its best expression in these issues of fine-tuning of parameters, the likelihood of a theory (or lack thereof), renormalisation conditions etc? Williams 2015 (section 3.2) has catalogued a variety of motivations for turning naturalness into a constraining principle for future physics:

1. **Fine-Tuning:** this is perhaps the most popular and straightforward way of presenting the problem. Having to set parameters equal to a very precise value looks like the kind of conspiracy that any sensible physics should avoid.
2. **Aesthetics:** Theories that rely on fundamental parameters, which are not of order 1 and/or widely differ in their size, may not violate some “rationalistic” criterion but they look “ugly” or undesirable and alternatives should be sought.
3. **Quadratic Divergences:** Results that contain quadratic sensitivity to some cut-off will exhibit a dramatic divergence when that cut-off is taken to infinity or large energy scales. In this sense, the cosmological constant problem belongs to the same family of future physics problems as the hierarchy problem.
4. **UV Sensitivity:** Since physics across scales should decouple, parameters at a low energy scale should not depend on degrees of freedom that only become manifest at higher energies.

Since we have already discussed in detail the possible formulations and drawbacks of fine-tuning arguments in the previous section, we will mostly concentrate on the remaining “sources” of motivation. As we are about to see, it is safe to assume that the problem of quadratic divergences can be reduced to fine-tuning or UV sensitivity. Fine-tuning itself can be either seen as exemplifying UV sensitivity or some aesthetic preference. Hence, when the dust settles, it all comes down to the dilemma: is naturalness grounded in a physical insight or some aesthetic preference?

Fine-tuning Intuitively, fine-tuning seems exactly the kind of problem that calls for an explanation. The challenge here, however, is to turn this intuition into a cogent argument. We

examined the main obstacles to achieving this in the previous section. Needless to say, even lacking a strong, formal argument, one can still adhere to their intuitive aversion to fine-tuning. After all, as is often the case with insights, their existence precedes their essence: even glimpses of the truth could be the way forward in the otherwise tenebrous ambiance of BSM physics. However, lacking a clear rationale to disclose the deeper issue lurking beneath the muddy waters of intuition, one will have a hard time persuading the (obstinate?) dissenter whose intuition will remain unfazed by such “intuition pumps”. Therefore, the intuitive basis for fine-tuning is not strong enough to take the hierarchy problem seriously.

Aesthetics Evidently, aesthetic judgments are quite hard to assess: they seem too subjective, taste-dependent, lacking a rational or, sometimes, any sort of foundation and frequently are just arbitrary. For some, Pollock’s “One: Number 31, 1950” might constitute the hallmark of the expressive power of abstract art while for others it is the very epitome of the things that have gone wrong with it. While criteria and deliberation can help resolve this impasse even in art theory, it certainly is *prima facie* hard to see how a scientific theory might be said to be aesthetically appealing, beautiful or elegant without adding too much hand-waving to the mix. Ultimately, such impressions will also need to be grounded in something more tangible to influence the way scientific theorising will be conducted.

Aesthetic considerations can enter the problematic in many ways, but here we will focus on two. The first, more practical and more directly relevant to the above is determining the amount of fine-tuning that is acceptable. Since, no physical law puts any independent constraints on this, it is left to the physicist, philosopher or whomever to define what *they* take to be undesirable and what not. This, however, has a casuistic spirit that had better be avoided when discussing guiding principles. A more abstract and generalisable characterisation of beauty in scientific practice is desirable. Perhaps the most conspicuous such characterisation is **simplicity**. Any theory feels better, more elegant, preferable when it is simple. Usually, this reflects its capacity to incorporate a multiplicity of disparate phenomena under a common framework constructed by a minimal set of basic principles. In this sense, Newton’s unification of stellar and earthly phenomena, the resolution of the tension between mechanics and astronomy and the postulation of small set of laws governing all physical processes is the exemplar of beautiful theories. QFT and the SM can also be said to be elegant in this sense since, using the mathematical framework of group theory, they have led to a spectacular classification of the

fundamental constituents of matter and their interactions.

Yet, this is far from satisfactory motivation in which to ground the perceived role of naturalness. First, it is still unclear whether the SM can truly be deemed an elegant theory. Actually, many scientists would find the number of unspecified parameters and its failure to incorporate gravity as ugly features. Needless to say, it is not at all self-evident why putting parameters by hand is such an ugly feature a theory might have – especially if most physical theories have had free parameters to adjust empirically. The more important challenge, therefore, is to turn this premature and fuzzy insight into a more substantive and concrete principle that determines theory choice. And, clearly, simplicity cannot be something along the lines of “favour theories using the simplest mathematical tools available”. This way the Dirac equation would have been a disaster compared to the Klein-Gordon equation, which involved a rather straightforward substitution of the relativistic energy relation into the Schrödinger equation. Still, one cannot dispute that Dirac’s derivation can be seen as an attempt to preserve as much of the simplicity of the original equation (retaining linearity) by generalising ψ -s to matrices.

An even more lucid illustration of the perils of beauty is that, given its unclear nature, beauty is – to echo Hossenfelder 2018 – highly prone to “leading physics astray”. We have already seen how Dirac’s fascination with large numbers led him to a completely erroneous conclusion about the gravitational constant. The underlying motivation in that case was exactly the idea that something deeper will be revealed by these arithmological coincidences. Like the gambler who takes the number on his friend’s shirt as a sign by the universe (or the Almighty) about what number to bet on at the roulette, the scientist is supposed to take the appealing (or appalling?) numerical ratios to be a sign towards some deeper truth about the world. There is some kind of purposefulness in nature that is or will be captured by our scientific theories. Kragh 1990, who wrote a very detailed biography of Dirac, claims that his preoccupation with beauty at the expense of more empirical, hypothetico-deductive or pragmatic methods resulted in a series of setbacks in his later work:

In Dirac’s scientific life, the mid-1930s marked a major line of division: all of his great discoveries were made before that period, and after 1935 he largely failed to produce physics of lasting value. It is not irrelevant to point out that the principle of mathematical beauty governed his thinking only during the later period.

Whether or not this bold assessment of Dirac’s work is accurate, it is safe to say that although beauty can be of use as a pragmatic tool, its epistemic value should not be exaggerated. Quite

unsurprisingly, if naturalness is construed as a purely aesthetic principle of a pragmatic flair, it must be accorded a significantly weaker status and dispensing with it will be a rather painless process.

A more down-to-earth conception of the role of beauty in physics with an emphasis on simplicity and the unificatory role of mathematics was given by Einstein (quoted in Howard and Giovanelli 2019):

...in my opinion, the right way exists, and that we are capable of finding it. Our experience hitherto justifies us in trusting that nature is the realization of the simplest that is mathematically conceivable. I am convinced that purely mathematical construction enables us to find those concepts and those lawlike connections between them that provide the key to the understanding of natural phenomena. Useful mathematical concepts may well be suggested by experience, but in no way can they be derived from it. Experience naturally remains the sole criterion of the usefulness of a mathematical construction for physics. But the actual creative principle lies in mathematics. Thus, in a certain sense, I take it to be true that pure thought can grasp the real, as the ancients had dreamed.

Einstein's more concrete understanding of simplicity takes flight from the sky-lands of nephelococcygia to address the thorny of theory construction. With mathematics serving as the vehicle for refining, redefining and expanding the available frameworks. Ironically for Einstein, the primary example of the successful application of this strategy is QFT for which showed little interest during his later life. The central role of symmetries (the generalisation from abelian to non-abelian groups to include more interactions, the representations of the Lorentz group, the spin-statistics connection etc) attest to the correctness of Einstein's insight.

Yet, even in this sensible approach to beauty, there are still two outstanding issues to be resolved. First of all, Einstein's unificatory methodology only covers a narrow aspect of theory transitions and can easily lead to a misleading conception of the function of theories. Indeed, our more situated analysis of the strategies employed in making physical theories map onto the world would probably be exiled to the lands of ugliness by Einstein's standards. The whole EFT-based theorising is replete with choices over the scales, the relevant interactions, the values of parameters and so on. It completely does away with the kind of fundamental theory talk that would be most congenial to Einstein's spirit in the above quote. In any case, it is hard to see how naturalness in any of the forms we discussed could fit into this methodological scheme. Little does it share with the vision of a systematic construction of physics on the basis of some creative

principles inherent in mathematics. As we are about to see, naturalness at best correlates with the decoupling of scales and its abandonment (as a form of obtaining a generalised principle across a wider range of phenomena) does not correspond to an obvious mathematical fact - in the manner the axioms of SR codified the information about Lorentz transformations, for example.

Quadratic Divergences Quadratic divergences for scalar fields like the Higgs field can still be eliminated using standard renormalisation techniques but they leave residual terms $\sim M_f^2$ that receive contributions from all fermions coupled with the field, even from scales well-above the examined EW energy scale. We have seen how critics of naturalness have precisely attacked it as an artifact of the regularisation scheme – a clear indication that they dismiss it as a mere aversion to quadratic divergences. But what is particularly problematic about quadratic as opposed to logarithmic divergences? In the end, we cannot help but fully agree with Williams 2015 in his verdict that “quadratic divergences cannot actually have been the central issue in the first place” (p. 88). The key to understanding the real problem with this type of divergences is their connection with parameters of relevant operators such as the Higgs mass and the cosmological constant.

UV Sensitivity The real cause of uneasiness with the existence of quadratic divergences is that they illustrate in the most dramatic fashion the sensitive dependence of low energy physics on high energies. Physicists sensitive to these conceptual issues (e.g. Craig 2017, Koren 2020) have been particularly heedful of this potential failure of decoupling and claim that this is precisely the physical meaning that should be attached to naturalness as a principle. This construal of naturalness, albeit less precise than more technical formulations, will be much harder to dismiss. Guaranteed plausibility comes at a price, however: much like in the interpretation of an aphorism, water down the point and all one is left with is a truism. Unsurprisingly, we need to turn to an assessment of this more “watered-down” version of naturalness and the implications of its failure for the effective framework.

3.5 Reaction 3: From Effective to Ineffective Theories?

Treating naturalness as synonym for decoupling has at least two advantages for its proponent. Evidently, this variant is not susceptible to the objections and misgivings of more technical for-

mulations. At the same time it reveals the profound physical intuition that remains for the most part shrouded in discussions of fine-tuning, beauty etc. The demand for decoupling of scales is a physically “thick” notion much like Lorentz symmetry, locality or energy conservation. Williams 2015 presents decoupling as *the* definition of naturalness codifying what he calls the “Central Dogma” of the EFT program:

Ultimately, I claim, the reason that failures of naturalness are problematic is that they violate a “central dogma” of the effective field theory approach: that phenomena at widely separated scales should decouple.

If this is right, then violations of naturalness have pretty damning repercussions, indeed. They indicate a breakdown or, at the very best, a restriction of the EFT program itself: the autonomy of scales might not be the universal property of nature we originally thought. Sometimes it just might be necessary to know the tiny details of microphysics to derive the correct macrobehaviour. Accordingly, the applicability of EFTs is perhaps limited to cases where a set of pre-conditions is satisfied. For example, these could be the assumptions needed for the Appelquist – Carrazone theorem to hold, which while reasonable, soon prove to be restrictive (e.g. Georgi 1993, sect. 2.5):

1. There is a separation of scales for the massive field case.
2. Decoupling holds between the low and high energies when the higher energy theory is perturbatively renormalisable.
3. The theory must employ a mass-dependent scheme because mass-independent schemes (such as minimal subtraction) will lead to incorrect logarithmic sums.

The decoupling theorem holds in QED and QCD but fails to apply to the weak sector the SM as “we cannot parametrize and match together effective theories by switching off fields of mass $M > \mu$ at a given scale μ ” (Jegerlehner 2015, p. 1181). Nonetheless, one usually does not expect this formal result to hold before attempting to construct an EFT - decoupling is frequently “vindicated” by producing accurate results when the matching is “put in by hand” (Georgi 1993, section 3.1., Manohar 2018, section 3.5). What Bain 2013 has dubbed “continuum EFTs” are precisely constructed not by integrating-out degrees of freedom (with Λ being a cut-off that splits the energy levels), but “forcing” the two theories to agree at a certain energy scale. This shows that the applicability of EFTs is not fully exhausted by the assumptions we find in the decoupling theorem.

The Stakes Let's grant that violations of naturalness imply that the autonomy of scales is violated as well. What are the repercussions? This depends, I think, on the scope of this failure:

- **RESTRICTED:** If decoupling is seen as encoding a particular feature of QFT, then its violation can be seen as a singularity / peculiarity of this specific framework that will be resolved through some re-conceptualisation or perhaps minor adjustments. Biting the bullet in this case amounts to an abandonment of the way of thinking about these issues that has become canonical in the course of 20th century particle physics. Physicists seem to understand naturalness in this more constrained manner.
- **UNRESTRICTED:** If decoupling is understood generically as involving the whole of physics, then the stakes will become significantly higher. Major methodological revisions will be needed; reductionism itself might come under heavy fire. This is the more dramatic flair added in a recent paper by Wallace 2019.

Ecce dilemma: barring any amendments, decoupling, in the more restricted, technical sense we come across in QFT fails or is generally undermined in physics eventually carrying reductionism down along with it. Before plunging into the depths of despair, however, I think that there are two potential ways out of this Scylla-Charybdis trap:

- It is not established that naturalness, with the prohibitions it introduces to the values of lower energy parameters, is to be equated with the the autonomy of scales requirement. The decoupling achieved in “manual” (bottom-up) versions of the matching procedure is empirically motivated and justified; while not grounded in some specific theoretical result (like the AC theorem for top-down version), it is vindicated by the predictive success this approach enjoys. In actual practice, this is the more useful technique since in most interesting cases the higher energy theory is unknown. Now, it is certainly the case that some theorem à la Appelquist-Carrazzone might exist to establish the necessary form of decoupling, but at least from the perspective of practice this is not needed to make use of the EFT framework.

⇒ **Suspicion:** Naturalness is a more stringent species of decoupling, i.e. one that puts constraints on the values of parameters. Theories (as evidenced by the SM despite the hierarchy problem) can still be predictive and fully functioning with some more restricted (the unnatural kind) form of autonomy at play.

- An even more interesting possibility is to “quarantine” this failure of decoupling to prevent its disastrous side-effects from “spreading” to other domains of physics. If such violation only occurs within the narrow confines of QFT, because, say, we have reasons to believe that it is in some respect unique or dissimilar to other theories, we can downgrade the problem to a peculiarity of QFT; perhaps all we need is some new conceptual developments internal to QFT to resolve it.

⇒ **Suspicion:** In lieu of a systematic collapse of the reductionist program, violation of naturalness only concerns QFT and indicates a possible restriction on the applicability of the effective framework.

We will examine each of the two possibilities in turn. Prompted by a recent paper by Wallace, we will start by granting the assumption that “naturalness = decoupling” and then attempt to deflate the problem by differentiating QFT from theories examined like statistical mechanics. We will then continue with the first possibility, which puts into question the very identification of naturalness and decoupling. If either of the points holds, pace Wallace, violations of naturalness do not have cataclysmic repercussions. If the stronger point about naturalness \neq decoupling holds, then EFTs themselves are not under severe pressure.

3.5.1 A Failure of Reductionism?

Wallace 2019 mounts his defence of naturalness suggesting that, as a principle, naturalness codifies a form of reasoning pervading all of physics. He argues that should naturalness be deemed irrevocably beyond repair, the whole edifice of physics must be taken to lie on pretty shaky (methodological) foundations:

This means that if naturalness really fails in high-energy physics, that failure undermines not just fine details of particle physics but the entire hard-won understanding we have of how physics describes systems at different levels and how those descriptions interrelate with one another. The apparent failure of naturalness is then a crisis at the heart of contemporary physics.

Put in pretty dramatic terms, indeed! Yet, who can blame the tone if, as Wallace claims, a failure of naturalness essentially amounts to a failure of reductionism, the central (methodological, ontological or you-name-it) dogma that has guided particle and most of 20th century physics? His argument is worth examining in detail as it rests on an interesting and seemingly compelling

analogy with considerations of typicality in statistical mechanics. Let us track his thought in some detail.

Wallace begins his argument by analogy with a discussion of explanations in the context of statistical mechanics. As is known, a statistical mechanical treatment of a given system begins with its phase space. Each point in this space represents the position and momenta of all particles making up the system at a specified time t and is called a microstate of the system. Knowledge of the microstate along with the laws governing dynamical evolution suffice to fix the microstate of the system at a later time. Tracking all these degrees of freedom is an exceptionally hard task, however. So, we have to resort to coarse-graining: partition the phase space into cells of aggregate microstates that satisfy some coarse-graining condition (e.g. velocities are close to some value). The Liouville measure, a smooth, uniform measure is then used to assign some probability to each cell. As a result, despite the huge number of microscopic degrees of freedom involved, we are in a position to make probabilistic claims about the evolution of macrostates and assign some probability to possible dynamical histories.

The key feature of the systems treated in statistical mechanics is what known as their “mixing property”. This property guarantees that the exact form of the measure adopted will be irrelevant for obtaining the right macroscopic behaviour. Nonetheless, some caution is advised: the chosen distributions should not be highly irregular, i.e. significantly deviate from the Liouville measure, lest one run the risk of obtaining highly erroneous conclusions²⁸. Naturalness enters this picture as the extra assumption required to eliminate these problematic distributions and derive the typical macroscopic behaviour of systems. The idea is to classify initial distributions as follows:

1. **Natural**: a probability distribution on the phase space of the system that leads to some stable classical macrodynamics; these distributions over the microstates consistently lead to the typical macroscopic behaviour
2. **Unnatural**, i.e. distributions that do not satisfy the above conditions and are further subdivided into:
 - (a) **Weakly Unnatural**: probability distributions that fail the “naturalness” requirement but, which, thanks to the mixing property of statistical systems, will also give

²⁸See chapter 4 in Myrvold 2021 for a nice illustration of this issue using a parabola gadget as a simple toy example.

rise to stable macrodynamics. In other words, any potentially distorting details are “washed out” by the laws of dynamical evolution.

- (b) **Strongly Unnatural:** probability distributions so pathological that the dynamics fails to wash out their peculiarities and will result in highly atypical macro-behaviour (like ice-cubes spontaneously forming out of water)²⁹.

Wallace’s next claim is that this description bears strong resemblance with EFTs in the QFT framework. Probability distributions here are range over the space of physical parameters such as masses, coupling constants etc. As we know, these parameters depend on the cut-off scale at which the EFT breaks down and their values change as the system energy varies. They are classified as irrelevant, relevant and marginal³⁰.

Wallace, then claims, that the analogue of the Liouville measure in statistical mechanics is a uniform measure over the space of these parameters. Any measure that is “tractably specifiable” in terms of the uniform measure is considered to be natural. Natural measures render unlikely the possibility that relevant parameters will take on very large values and thus dominate the low-energy physics. The low-energy physics will be natural if it arises from a natural or weakly-unnatural distribution over parameters at the high-energy scale [we can also define “weakly” and “strongly” unnatural large distance physics in a manner analogous to the Stat Mech case: some unnatural distributions will give rise to natural “macro”-dynamics]

Therefore, for Wallace, naturalness is strongly tied to our picture of emergence in physics. The naturally emergent macrodynamics is compatible with (almost) any possible microdynamics. We need some naturalness assumption to derive macrobehaviour from microbehaviour:

$$\boxed{\text{Micro-dynamics} + \text{Naturalness} \longrightarrow \text{Macro-dynamics}}$$

Conversely, an unnaturally emergent macrodynamics would require an atypical (or unnatural) initial distribution over the microscopic (or higher energy) level. In a world in which macrodynamics was unnaturally emergent, most of our physical knowledge would be about intricate

²⁹Essentially, the initial distributions trap the evolution of the system in a “spiral of (atypical) doom”. The dynamics cannot pull the system out of the “event horizon” of the irregular state it has found itself.

³⁰To help the reader skip some page turning, here is a brief reminder of what these terms mean:

1. irrelevant parameters: their value decreases as the system tends to lower energy levels; they do not influence IR infrared physics
2. relevant parameters: their value increases as we go to the IR dominating low energy physics
3. marginal parameters: their values remain approximately constant throughout the RG flow; their behaviour needs to be examined more closely to see how they influence IR physics

details of the initial conditions and precise value of the parameters. Contrary to what we expect from our current understanding of areas like statistical mechanics, chemistry, astrophysics and so on, the details at the microscopic level could not be ignored or coarse-grained without losing the ability to produce accurate predictions for the macroscopic systems at hand. Wallace continues with a discussion of the possible justifications for naturalness, but we will leave this aside for the moment.

How compelling is the above argument for naturalness? While nicely underlining the affinities between different areas of physics, its persuasiveness ultimately rests on how strong the analogy between QFT and CSM is. Our contention is that Wallace's argument is vulnerable across multiple fronts:

- Assumptions: Wallace makes some implicit and explicit assumptions, the truth of which is far from uncontested:
 1. Analogy 1: the EFT techniques used in QFT constitute another form of coarse-graining very similar to the coarse-graining procedure one encounters in statistical mechanics
———— We find that this is disputed by recent work by Fraser (and Koberinski) on analogical reasoning in QFT.
 2. Analogy 2: the uniform measure over parameters in QFT is the analogue of the Liouville measure in statistical mechanics
———— Apart from considerations over the meaning of probabilities discussed in the previous section, there is a clear dissimilarity in that the Liouville measure is picked out by the dynamics³¹
 3. Presupposition: the “naturalness” strategy as described in StM actually works in that context
———— Problems with the reduction of thermodynamics and the arrow of time seems to indicate that this is also a controversial statement.
- Conclusion: Wallace ignores or downplays the viability of some alternatives out of the conundrum he has led us into:

³¹The Liouville measure is accorded special status in StM because it remains invariant under the evolution of the Hamiltonian describing the system.

1. Option 1: Restrict the domain of applicability of EFTs
 —→ “Effective” in the more informal sense
2. Option 2: Reconsider the notion of reductionism we are working with
 —→ Compositionality without derivability.
3. Option 3: Revise the notion of decoupling
 —→ Historical precedent: rejection of separability

We will investigate each of these possibilities to argue that they provide sufficient ground to reject Wallace’s punchline:

Without Naturalness, on the other hand, the connections between physics at different levels are severed and we lose any ability to understand inter-level relations. This means that the longstanding failure of naturalness in cosmology, and the more recent evidence for Naturalness violation in particle physics, have ramifications far beyond those specific and esoteric fields. Naturalness failure here undermines arguments for Naturalness anywhere, and calls out urgently for understanding.

The style of reasoning across the domains examined might be very similar, but this misleadingly conceals some important disanalogies between them to artificially magnify the repercussions of specific problems at the frontiers of fundamental physics.

Assumptions Returning to our “Decoupler’s Dilemma” we start with the assumption that naturalness is a global, i.e. “ubiquitous” principle applying everywhere in physics. In this case, there are two possible ways out depending on whether one grants the *naturalness = decoupling* [$N = D$] hypothesis. If $N = D$ is granted, then to avoid having failures of naturalness in QFT leak into other areas of physics such statistical mechanics, one can deny that the the two are “sufficiently” analogous. We will base this argument on recent work by Fraser (D. Fraser 2012, D. Fraser 2020) which specifically examines the nature and role of analogies in QFT. Fraser’s key insight is that these are usually analogies of a thinner, formal or mathematical kind rather than the thicker, physical kind. Fraser has examined analogical reasoning with respect to spontaneous symmetry breaking in QFT and CSM, the Higgs mechanism and condensed matter physics (in her 2016 paper with Koberinski) as well as renormalisation in QFT and CSM. As these arguments are structurally similar, we will focus here on the latter, which is more pertinent to our discussion. An alternative is to deny that $N = D$ and hence

disentangle naturalness from decoupling by refining both concepts perhaps in the manner of Bain 2019 and Franklin 2020. We will devote a section to each of these possibilities and then return to more revisionary alternatives (targeted at Wallace's conclusions) next.

3.5.2 Quantum and Classical Analogies

A. Analogies Let's start with some terminological distinctions. Suppose we have two events, objects, domains or, as is the case here, theories A and B, which share many features and A is more thoroughly understood or easier to work with than B. Analogical reasoning uses this affinity to make inferences about B on the basis of what is known about A. This is the case with the study of the causes of wars in history, for example. We usually find that there are certain deep commonalities in the structure of societies, the needs of peoples, the availability of resources etc across all historical eras. We use this observation as the ground on which to "transfer" expectations about one historical incident to another; similar causes will be at work in well-studied cases like WW2 and earlier events such as the Seven Years' War. Somewhat analogously it is not uncommon for scientists to use structural similarities between fields or models of physics to transfer results from one field to the other.

Fraser adopts Hesse's useful categorisation of analogies as employed in physics first distinguishing two types of relations:

- Horizontal: mappings from the domain of one theory to the other (e.g. the shape of tennis ball and that of the sun)
- Vertical: mappings between objects in the domain of the same theory (e.g. the oscillation of a surface and the production of sounds)

Vertical relations are important for preserving the physical content of the theory. Depending on the type of relations - correspondence mappings between any two theories one can define three types of analogies:

1. **Formal Analogies:** The "thinnest" instance of analogical reasoning is when there is only a mapping between interpretations of the same uninterpreted theory - like applying the same equation in different fields.
2. **Physical Analogies:** A "thicker" analogy exists if the horizontal relations reflect some deeper (physical) similarities between the two domains.

3. **Material Analogies:** If among these similarities we find causal relations of the same kind we get a stronger subspecies of physical analogies.

We can now see how this classification works out in some example. A clear physical analogy would be the application of the mechanics of waves to sound, seismic or surface waves. One finds that: a) the descriptions share a common mathematical structure (wave equation) and b) a similar physical basis (propagation through some medium). By contrast, the case of electromagnetic waves is not a physical analogy as the physical basis and mechanism of propagation (in the sense of molecular vibrations and interactions making up the wave) is missing. Therefore, it should not surprise us to find some important physical dissimilarities between the two.

QFT \sim CSM D. Fraser 2020 presents a very detailed case study of Wilson’s groundbreaking work on renormalisation techniques (K. G. Wilson 1975, K. G. Wilson and Kogut 1974) as they are applied to quantum fields and systems studied by (classical) statistical mechanics. Wilson’s achievement was an impressive feat of analogical reasoning. His ingenious insight was to introduce a correspondence between some mathematical structures of both theories and then exploit methods typically used in the context of CSM to systems studied in QFT. Fraser breaks this down in three conditions:

1. (Identity): Wilson identified correlation functions $\Gamma(x, y)$ in CSM, which encode (spin)³² correlations between fluctuations at point x and point y in space with a Wick-rotated propagator $D(x, y)$ in QFT, which (typically) describes the probability for a particle (an excitation of the field) to propagate from spacetime point x to y :

$$\Gamma(\mathbf{x}, \mathbf{y}) = \zeta^2 D(x, y) \tag{3.56}$$

where ζ some scale factor. Note that x, y also have a temporal coordinate, which, per Wick-rotation, is now $\tau \rightarrow it$. The propagator can thus be written, to more closely resemble Fraser’s notation, $D_{\Delta x}(-i\tau)$.

2. (Constraint): A correspondence between the momentum cut-off Λ in QFT and the corre-

³²This depends on the system as well as the kind of correlations one examines. A classical example found in most introductory texts in the literature (e.g. McComb & Tauber 2005) is some n -dimensional (usually two- or three-) configuration of spins $s(x)$ and their alignment (up or down) and the way an external magnetic field influences them.

lation length of the system ξ is established:

$$\Lambda = \mu_R \xi \tag{3.57}$$

with μ_R the renormalised mass for the quantum field. This constraint guarantees that the propagators will be well-behaved as the lattice spacing is sent to 0.

3. (Limit): In the end, one sends the momentum cut-off to infinity, $\Lambda \rightarrow \infty$, in order to return to a continuum theory. The constraint keeps the propagators well-defined throughout this process; μ_R takes on a fixed value furnished by experiment.

Fraser then goes on to examine in detail the specifics of each of these steps to demonstrate that they can at best be seen as establishing formal, not physical, analogies. We summarise the main points here for quick reference:

Table 3.2: Wilsonian QFT & Classical Statistical Mechanics

Formal Analogies	Physical Disanalogies
<p>IDENTITY: The mathematical quantities in both theories fulfil the same mathematical role (e.g. the spin field $s(\mathbf{x})$ and the scalar field $\phi(x)$ are the dependent variables and the correlation functions and propagators are expectation values of products of fields)</p>	<p>The correspondence between Γ and D is a mapping from space to spacetime. Time has a significantly different role in theories (intimately related to causality). Therefore, there is no common causal structure between the two cases.</p>
<p>CONSTRAINT: i and μ_R^{-1} play an analogous role in the correlation functions: $\Gamma_{\mathbf{x}} \propto e^{\frac{\mathbf{x}}{\xi}}$ $D(-ix\alpha) \propto e^{\mu_R \alpha}$ with α determining the lattice spacing.</p>	<p>Contrary to xi, μ_R does not represent any spatial quantity (minimum size at which the system loses its qualitative character) and $\alpha = \frac{1}{\mu\xi}$, the lattice spacing in QFT is conventional (introduced to fix the units for the physical quantities correctly) and is finally taken to 0. This is not the case for ξ in CSM where no continuous limit is required.</p>

LIMIT: One exploits the previous analogies to take the continuum limit of QFTs, $\Lambda_0 \rightarrow \infty$. One goes through the space of theories coarse-graining and re-adjusting the parameters so that they converge towards the theory with the right (physical) renormalised parameters.

Parameters in CSM such T_0 are experimentally accessible in, but this is definitely not the case for bare parameters likes m_0 in QFT. Taking $T \rightarrow T_c$ represents the physical process of temperature change, but no physical process is represented when varying μ_0 . Also, while the temperature is causally contributing to $\xi \rightarrow \infty$, this is not the case for μ and Λ_0 .

What is the relevance of this discussion for naturalness and the hierarchy problem? This becomes clear when we remind ourselves of the fact that³³ “much of our intuition for fine-tuning and naturalness comes from condensed matter physics” (Schwartz 2014, p. 411). After Wilson’s work in the 1970s it has become become standard practice to inform our understanding of RG equations in QFT by the mathematical analogue we find in CSM, frequently oblivious to their underlying differences (see Fisher’s contribution and comments on Weinberg’s talk in Cao 1999). This has solidified the expectation that results, insight or problems from one domain carry over the other. It is unsurprising, then, to come across claims like the following passage from Schwartz 2014 (p. 410)

Suppose the theory were finite, for example if it were UV completed into string theory, or more simply if it were the effective description of some condensed matter system (in which case Λ might represent some parameter of the microscopic description, such as an inverse atomic spacing). Then the bare mass m and cutoff Λ would be physical.

Apart from the more technical issues presented by Fisher 1998 (p. 89), Fraser’s and Koberinski’s work on analogies between the two fields is further indication not only that the physical basis of the RG techniques used in QFT or SM is dissimilar, but also that the analogy has overshoot its target. Several open problems add momentum to this suspicion. On the one hand, there is the dubious status of cut-off QFT:

³³The same spirit is arguably echoed by Peskin and Schroeder 1995 when they write “... in order to end up with the desired value of m^2 at low momentum, we must imagine that the value of m^2 in the original Lagrangian has been adjusted very delicately. This adjustment has a natural interpretation in a magnetic system as the need to sensitively adjust the temperature to be very close to the critical point. However, it seems quite artificial when applied to the quantum field theory of elementary particles, which purports to be a fundamental theory of Nature. [...] Perhaps this is the reason why there seem to be no elementary scalar fields in Nature.” (p. 406)

1. The status of the cut-off in QFT is debatable. Contrary to the clear physical significance of the cut-off in CSM, the existence of alternative regulators in QFT shows that the cut-off may be more akin to a formal tool.
2. This point is reinforced by the continuum version of RG, in which the scale μ introduced does not involve a separation of degrees of freedom into higher and lower ones³⁴.
3. Even, more dramatically, the possibility of alternative formulations of QFT leaves open the possibility of a non-effective conceptualisation of QFT.

Nevertheless, one does not have to go this far. Some further disanalogies may be noted:

1. There is a leap involved when shifting from distributions in StM to distributions in QFT. In the former, we characterise the kind of micro-configurations that typically obtain, something we have plentiful empirical evidence for. By contrast, probabilities over parameter space in QFT refer to variables we cannot sample. The best reply here would be to shift to Bayesianism, but, as we saw in the previous section, the success of the manoeuvre is at best dubious.
2. An obvious, perhaps naive, disanalogy to note is that the RG in StM has a clear interpretation as a form of coarse-graining whereas RG techniques in QFT can be come in the form of coarse-graining or integrating-out. These two may come apart at times.
3. The above points naturally lead us to what is perhaps the fundamental difference between naturalness reasoning in StM and QFT. For StM naturalness is akin to a transcendental guarantor that the microdynamics will lead to the expected macrodynamics. For QFT the most cogent position we can take is that along the lines of Bayesian inference: what values of the higher energy parameters render more likely the lower energy theory³⁵.

I take it that all these cautionary remarks signal that it is at least premature to extrapolate a failure of reductionism in all of physics from this apparent failure of decoupling in QFT. More generally, since the physical analogy between CSM and QFT is at best debatable, drawing conclusions similar to Wallace's is a precarious business. Consequently, the apparent QFT

³⁴It has recently been claimed by Rivat 2019 that this renders the Wilsonian, cut-off approach more “physically transparent” and therefore preferable as the conceptual basis for understanding the RG in QFT. Besides contradicting the current practice of working with the continuum brand of RG, switching to the cut-off version would be extremely complicated for gauge theories (e.g. Schwartz 2014, sect. 23.6, Manohar 2018, 3.5).

³⁵And recall, the problem here, as emphasised in our brief recounting of BSM physics, is that we have been able to find experimental support for these “more likely” (for the most part supersymmetric) models.

failure of naturalness should not be deemed as begetting a profound revision to the way we globally understand physics. Note also the peculiar nature of reasoning involved: we first hypothesise that a certain analogy between domains X and Y holds, next we find out that there is something problematic about domain Y, but (here is the peculiar twist!) instead of taking this as contradicting evidence, we persist with the analogy and are now puzzled about X as well. A more reasonable, cautious stance would be to either forgo the analogy or restrict its scope. This approach is more sensitive to the fine points of departure between theoretical frameworks that at first glance appear very similar.

B. Presupposition Of course, an implicit, though definitely not controversial, presupposition for Wallace’s analogical argument is that the naturalness story he offers is accurate even for statistical mechanics. As is well-known, a long standing problem in the foundations of statistical mechanics is whether and how thermodynamics, entropy increase and macroscopic irreversibility is derivable or at the very least compatible with the time invariant microscopic dynamics of classical mechanics. If anyone follows the far more radical path of disputing the kind of reasoning laid out by Wallace even for statistical mechanics, then the whole argument does not even take off the ground. Of course, there is a price to pay when going against what has been the norm in the field of statistical mechanics. One would presumably need to address questions like:

- What is the status of the 2nd law of thermodynamics?
- If 2nd law is an empirical generalisation: What is the explanation for almost ubiquitous time irreversibility in macroscopic (sub)systems?
- If it is a contingent fact stemming from the dynamics and specific initial conditions: How come initial conditions play such an essential role in deriving the correct macroscopic behaviour?

Perhaps an account that offers answers to the above is a plausible alternative to the “micro to macro formula” defended by Wallace. For an attempt to formulate such an alternative, see Zelko forthcoming. For our purposes, we will remain neutral with respect to the plausibility of such a project and focus on the QFT case.

3.5.3 Naturalness \subset Autonomy

Even if we grant that Wallace's argument works as described in StM and accept the force of the analogy with QFT, there is still room for manoeuvre for those reluctant or unwilling to abandon the possibility of a more substantive, physical, affinity between the mathematics of QFT and CSM. In fact, even the effective realist or any sympathiser of the view we examined in previous chapters will wish to go beyond noting the mere disanalogy between the two fields as a response to the failure of naturalness. This is because in the absence of an alternative to naturalness it might be impossible to support the basic ontological picture of semi-autonomous physical domains. Such an alternative route involves a deflation of decoupling - one that is detached from naturalness. To start walking down this path, let us surmise, pace Williams 2015, that naturalness does not really reflect the "central dogma of EFTs", i.e. the decoupling of scales, but rather imposes a much stronger requirement on a theory's parameters. Very recent work on naturalness and EFTs has been aimed at explaining the origin of the effectiveness of EFTs without an appeal to the value of parameters or similar principle. The main strategies can be broken down into two, not necessarily unrelated, classes:

- (Str1) introduce an alternative conception of decoupling
- (Str2) switch to a more informal, thinner notion of EFTs that remains unaffected by violations of naturalness.

(Str1) Franklin 2020 is a clear representative of the (Str1) camp; his paper is an attempt to account for the success of the EFT framework in a manner that does not hinge upon the validity of naturalness. Essentially, this means that scales do decouple from one another even when parameters are obtained by adding contributions from higher energy scales. To make this point more precise, he offers the following useful distinction between two senses of autonomy:

- ⇒ Autonomy_{ms} (autonomy from microstates): invariance of low-energy theory dynamics under changes in the state of the high-energy theory
- ⇒ Autonomy_{ml} (autonomy from microlaws): invariance of low-energy theory dynamics under changes in the parameters or laws of the high-energy theory; alternatively put: invariance with respect to changes that lead to different possible worlds!

A renormalisable theory is clearly autonomous_{ms} because the effects of high-energy scales can be absorbed into the redefined parameters. This means that these higher order processes barely

contribute to low-energy descriptions, which, as consequence, are largely insensitive to what is taking place at really small distances. Yet, this is no warranty that autonomous_{ms} theories will not violate the second requirement. The parameters of a low energy theory may still be very sensitive to those of high energy theory and any alterations to the latter will lead widely divergent results for the former. Franklin’s key point, however, is that even in when a theory is not autonomous_{ml}, a form of “effectiveness” is still at play. Thanks to renormalisation, one can set aside details about the structure of short distance physics and still generate empirically accurate predictions. Franklin helps himself to a notion of renormalisability that he dubs “effective renormalizability”: an effectively renormalisable theory may fail to be renormalisable simpliciter, but it will work just like a renormalisable theory up to a certain energy scale. This essentially means that up to the given energy scale, one will only need to specify a finite number of counterterms to produce finite results. This notion of effective renormalisability became conceivable and attainable with the advent of EFTs; as we’ve remarked already non-renormalisable were henceforth treated on par with renormalisable theories (up to a scale).

With this distinction at hand it is clear that natural EFTs are those theories that are both autonomous_{ms} and *autnomous*_{ml}, whereas unnatural EFTs only satisfy autonomy_{ms}. For this reason, if one refrains from equating decoupling with autonomy_{ml} and opts for autonomy_{ms} instead, then naturalness is not the principle explaining decoupling. This role should rather be attributed to effective renormalisability, which alone suffices to guarantee decoupling as autonomy_{ms} and hence account for the predictiveness of EFTs. In this vein, Franklin’s view echoes a relevant remark by Hossenfelder 2019 – emphasis added:

The change of parameters in the UV, which is done to quantify technical naturalness, is not a process that is physically possible. [...] At a given energy, these parameters have some specific values. We can’t change these values because that would amount to changing the laws of nature.

[...] *Decoupling is necessary to use effective field theory and is hence an assumption that underlies the whole framework of the renormalization group already.* This means the UV physics decouples whenever effective field theory can be used, regardless of whether or not the theory is natural.

Failure of naturalness does not imply failure of autonomy of scales tout court: we can still be largely or fully ignorant about high-energy physics and still obtain accurate predictions at the scales currently accessible. To add to the plausibility of this account we only need to remind

ourselves of our current predicament: we have been able to use the SM of particle physics to make predictions about the microscopic world even though a more fundamental theory is yet to appear on our radar screen. For all the shrouds of mystery engulfing it, the hierarchy problem has not come to haunt us yet in a way that hurts.

A very similar view is supported by Bain 2019 (p. 907-8) who distinguished between two forms of decoupling:

- **Heuristic decoupling**, as he calls it, is plainly the removal of high energy degrees of freedom and their inclusion in the redefinition of a theory's parameters. This can be performed by hand as in the bottom-up matching procedure.
- **Precise decoupling**, on the other hand, is a much more rigorous relationship holding between two scales – satisfying perhaps the preconditions for some more formal result like the AC decoupling theorem.

Bain's suggestion is that continuum approaches to EFTs, which recall are by far the most widely employed, will satisfy heuristic decoupling (as this is, at worst, put in there by hand), but only the Wilsonian approach will satisfy an even further more precise notion of decoupling (by integrating out higher energy degrees of freedom). Naturalness, in this sense, is an evidently stronger than needed requirement about the sensitivity of parameters on higher energies. It is "stronger than needed" in that the possibility of writing an EFT does not depend on whether it holds true or not. This is because in continuum approaches the parameters are always set at the physical scale of the problem and no assumptions have to be made about their value at higher energies³⁶.

(Str2) A proposal along the lines of (Str2) has been made by Crowther 2016. Seeking to establish a more neutral ground against those who think that EFTs have led to a radical form of anti-foundationalism (our "enthusiasts" of a chapter ago) and those who find the idea of a "final theory" inescapable, Crowther 2016 (sect. 3.8) suggests that we should adopt an "effective", that is more pragmatic and non-committal approach to the EFT framework itself. To this purpose she also puts forward her own contrast (p. 88, emphasis mine):

³⁶For the sake of completeness: Bain goes on to discuss the connection between naturalness and emergence to conclude that while not an appropriate principle for theory choice, it might still be worth examining from an ontological perspective as underwriting a particular way that scales in the world relate to one another. This latter point will not concern us here.

...distinguishing between different types of EFTs. Firstly, there are EFTs that are constructed from theories where the framework can be systematically applied—we might call these **formal EFTs**. In these cases, the low-energy degrees of freedom can be readily identified, and constructing an EFT is relatively straightforward, using the directions supplied by the framework. Secondly, there are those cases where we lack a formal way of identifying the appropriate low-energy degrees of freedom directly from the high-energy theory, and have to do a lot of work “by hand”, utilising other methods and data—call these **informal EFTs**.

Thus, one can still treat theories in accordance with the instructions of the EFT framework, while imposing matching conditions by hand and verifying its validity in practice when the transition from high-energy degrees of freedom to lower physics in a “smooth”, “derivational” manner is unavailable. Physicists are not alien to this approach. For example, Alvarez-Gaumé and Vázquez-Mozo 2011 write (p. 259):

...in the MS subtraction scheme, or any other mass-independent scheme, the decoupling of particles as we run from high to low energies has to be implemented by hand, integrating out the field that become heavy as we lower the energy. Thus, every time a particle threshold is found, the corresponding field has to be integrated out and the appropriate matching conditions on the low energy field theory imposed. Proceeding systematically in this way, we guarantee the correct decoupling of the heavy species while retaining the computational advantages of a mass independent scheme.

This renders EFT techniques applicable even when decoupling is not known to hold in advance. Decoupling is thus transformed from an a prioristic constraint on theorising to an empirically grounded principle whose scope and validity is tested whenever an EFT is applied.

Comment At this point, one might worry that our conception of EFTs or decoupling has become too diluted. With the informal take on EFTs, the dissenting voice continues, we have strayed too far away from the truer physical meaning of EFT techniques as eliminating high energy degrees of freedom. There are reasons to remain unfazed by this kind of worry, however. We just need to be reminded of them.

First of all, let us note that failure of, call it naturalness or precise decoupling or autonomy_{ml}, does not render QFT or the SM any less predictive. True, it might cast doubt on the way dimensional analysis arguments have been employed or might favour renormalisation schemes

that require some fixing of parameters, but this did not undermine the scope or validity of QFT as a framework for the description of small distance phenomena. But, wasn't this precisely the point, one might wonder? Isn't it supposed to be a violation of our reductionist inclinations that some parameters need to be set in a top-down manner? There are two replies to this. First, we need to remind ourselves of the fact that theories ever since the time of Newton have required parameter input of an empirical origin in order to be predictive. The demand for a derivation of parameters from some dynamical mechanism is an additional, stronger demand stemming from a semi-rationalistic, ultra-reductionistic demand of explaining every contingency from a set of basic assumptions.

Second, even when we translate the problem in the language of RG trajectories and some sensitive dependence on the initial conditions (here: the bare or fundamental parameters), this is not without precedent. Chaotic systems have long been known to exhibit precisely this peculiar behaviour: the emergence of macrophenomena such as hurricanes, pandemics or economic recessions sensitively depend on the micro-conditions set at the initial stages of evolution of these systems. Now, this feature of chaotic systems appears to have put traditional accounts of explanation and reductionism (by implying some form of holism – see Bishop 2017, sect. 5) under pressure, it does not entail that a full abandonment of reductionism is ante portas. Perhaps what the deflated conception of decoupling and EFTs is hinting at is a reassessment of what is meant to come out of the reductionist demand. If reductionism is meant to safeguard realism against the discovery of more fundamental theories or as a methodological tool for theory construction, it might be possible to avoid the hassle while retaining the benefits by a slight relaxation of our philosophical views. We will return to this in chapter 5. On the other hand, if reductionism is meant to preserve some ontological import, in the sense of compositionality, then violations of naturalness have nothing to say against. There is no tension with believing that the SM gives an accurate description of the fundamental building blocks of the universe even if some of the properties require careful adjustment.

Before closing this discussion, however, let me add one further consideration into the mix. The naturalness as autonomy of scales is indeed a more cogent understanding of this requirement, the common sediment left over by (William's) distilling process of multiple definitions of the term. But at the same time, this weakening or abstracting of the principle renders it less useful in theory construction. Aversion to autonomy violations only underline the negative claim that some form of parameter dependencies should be absent, but do not specify the

exact kind of feature the successor theory should possess. To make this clearer, contrast this “thinner” notion of naturalness with the “thicker” notion of technical naturalness. In this case, the principle is targeting a specific feature of the theory and indicates a clear way forward - looking out for a new symmetry. I think that this fuzziness places further credence to the idea that our notion of naturalness needs further refinement and is not to be equated with decoupling simpliciter.

3.5.4 Restrict Domain of Applicability

Yet another way out of the conundrum is to start by accepting that naturalness is indeed violated, naturalness does indeed codify some indispensable notion of decoupling and statistical mechanics indeed employs a similar principle when it goes from micro to macro. The only possibility left, then, apart from the one examined right afterwards, is to restrict the scope of the EFT framework. In this way, the culprit is not our assumption that EFTs are characterised by some notion of naturalness (thus granting Wallace’s point), but our implicit assumption that EFTs need to apply universally to all phenomena and across all energy scales. Why would such a move, that essentially takes QFT (if we assume that Wallace is right about the analogy) to constitute a singular case, be acceptable at all? There are, I think, two separate issues to examine here. One concerns the justificatory ground for considering naturalness as an all-encompassing principle. The other concerns the precise meaning of “effective theory”. Let us look at each one in turn.

Start with the grounds for accepting naturalness as having universal applicability. There are two ways to support this claim: one is as an empirical generalisation over the instances of successful application of the EFT framework, the other as a sort of transcendental precondition for doing physics. While Wallace’s argument is structured in a way reminiscent of the former, the force of the conclusion and the dramatic tone bring it closer to the latter reading. I take it that if naturalness has the form of an empirical generalisation, its abandonment or, even more innocuously in our case, its restriction, while unpleasant, frustrating or undesirable is a pill we will all eventually need to swallow in light of empirical evidence. However, a more transcendental understanding of naturalness will indeed necessitate cataclysmic changes in our very conception of physics! Despite these dramatic overtones, this has also not been without precedence. (Metaphysical)³⁷ principles such as that of sufficient reason, determinism, conti-

³⁷Here the term should be understood as more in line with Kant’s project of deriving the basic assumptions that

nity or separability were initially put forward as objections to hypothesis such as light quanta, probabilistic evolution or wave-function collapse³⁸ but were ultimately rejected or restricted thanks to the immense empirical success of incompatible physical hypotheses.

The next point concerns a disambiguation in the meaning of “effective theory”. Again, Crowther’s distinction between formal and informal EFTs will prove very useful. Frequently, one comes across claims of the form: “all physical theories are effective theories”. Such claims border on the level of terminological cliché by neglecting the fact that all the term “effective” refers to in this case is that theories have restricted domains of applicability. Biology, chemistry, sociology etc can only be effective theories in this sense. But they do not share the more formal apparatus we find in QFT with decoupling in the stricter sense of the AC theorem, the RG transformations and the behaviour of operators across energy scales and so on. Wallace makes his case using StM as his example, where the mathematical apparatus is analogous to that of QFT. However, this conceals the fact that these formal devices may not be as ubiquitous as claimed and the disastrous consequences follow only if we cannot take QFT as unique in some sense. This is certainly disallowed if we construe “effective theories” in the more permissive informal, as opposed to the more exclusive formal, sense. Evidently, conflating the two is not an innocent move to make.

Therefore, with the “restriction” move we are essentially claiming that the formal techniques used in EFTs might not apply in the case of the Higgs mass. The obvious weakness is that unless there is an explanation, some physical reason as to why the Higgs field is such an anomaly this response will be seen as a desperate ad hoc manoeuvre. Indeed, similar restrictions for EFTs will be presented in the following chapter when we examine another violation of naturalness – this time for the cosmological constant. However, contrary to that case, the physical insight in the case of EFT failure for the Higgs mass is not transparent. Whereas for cosmology (and gravity) one can identify specific issues (such as the form of background structure) that create complications for EFT techniques, nothing of that sort seems *prima facie* relevant for the Higgs case: violation of naturalness in this context seems an entirely internal issue of the standard model.

Nonetheless, this is not entirely true. We will discuss in more detail the pre-conditions required for the applicability of the EFT framework in the following chapter, but for now it is

render our scientific knowledge possible.

³⁸One would need to be a bit more circumspect about this particular claim here, but it is meant a label-term for the debate between Einstein and Bohr about the plausibility of the latter’s explanation of the EPR effect. Bohr’s view might be unsatisfactory on several other grounds but separability, as it turned out, was not one of them.

worth emphasising the possibility that the hierarchy problem has to do with a mixing of UV/IR physics. While prima facie this appears as a most bizarre idea, especially in light of the neat separation of scales we encounter in almost all phenomena we encounter, developments in the field of string theory and non-commutative field theory provide motivation for a wild violation of EFT expectations. UV/IR mixing in string theory has been studied for closed strings, but it is always present in a non-commutative field theory. The problem is that both are typically associated with gravitational effects and thus do not help us discern what is distinctive about the hierarchy problem. To this effect, it is worth discussing a model presented by Craig and Koren 2020 for ϕ^4 non-commutative field theory without gravity.

The basic assumption is to introduce non-commutative relations between position operators:

$$[x_i, x_j] = \theta_{\mu\nu} \quad (3.58)$$

with $\theta_{\mu\nu}$ being the so-called non-commutativity tensor. We are then led to a modified uncertainty relation:

$$\Delta\hat{x}_\mu\Delta\hat{x}_\nu \geq \frac{|\theta_{\mu\nu}|}{2} \quad (3.59)$$

which already shows that whenever one tries to create some wavepacket of some length in one direction, it will be forced to be elongated in the other to compensate. This signals the existence of a mixing or an interconnection between IR and UV modes.

One then introduces a new field product, the so-called Groenewold-Moyal product so that one can work easily in terms of fields whose coordinates are functions of commuting coordinates:

$$f(x) \star g(x) = f(x) \exp\left(\frac{i}{2}\theta_{\mu\nu}\partial_y^\mu\partial_z^\nu\right) f(y)g(z) \Big|_{y=z=x} \quad (3.60)$$

This new product allows us to define non-commutative actions and, which will be the basis for NCFT, such as the following for ϕ^4 theory:

$$S = \int d^4x \left(\frac{1}{2}\partial_\mu\phi\partial^\mu\phi + \frac{1}{2}m^2\phi^2 + \frac{g^2}{4!}\phi \star \phi \star \phi \star \phi \right) \quad (3.61)$$

The key difference for our purposes is that this action will lead to loop diagrams that planar and non-planar diagrams (reproduced from Craig and Koren 2020):

If the second, nonplanar diagram is evaluated it picks up a phase $e^{ik^\mu\theta_{\mu\nu}p^\nu}$, which mixes internal and external momenta with the promise of eliminating the UV divergence. Most intriguingly,

$$-\Gamma_1^{(2)} = \text{---} \overset{k}{\bigcirc} \text{---} + \text{---} \overset{k}{\bigcirc} \text{---}$$

Figure 3.6: Non-commutative Field Theory Loop Corrections

ingly, though, the non-commutativity of IR and UV limits implies that the standard (Wilsonian) techniques have to fail. This is because when we try to write down our EFT Lagrangian with parameters as functions of scale $Z(\Lambda), m(\Lambda), g(\Lambda)$ (the running parameters) we simply cannot include the term of the extra pole at lower energies since this only appears as a correction when take the $\Lambda \rightarrow \infty$ limit. In other words, physics at lower energies cannot be described by theories of the same form³⁹!

Even though toy models such as the above cannot still be said to solve the hierarchy problem (with another reason being that the new field’s propagation does not respect Lorentz invariance) they are nevertheless significant in showing a way in which physics can still function even when EFT expectations are violated. The fact that gravity has not entered the picture points to an internal insight about the structure of EFTs and the SM. Further developments along this path, especially when gravity will be finally taken into account in BSM physics, could prove that far from being all-pervading, the EFT framework faces its own natural limitations and should only be taken to apply under specific conditions. The hierarchy problem could in this case be a sign that we have reached one border of EFTland.

3.5.5 Beyond a Reductionistic Physics?

The much more radical possibility is to accept Wallace’s argument and lay blame on the reductionist agenda that has dominated physics throughout the 20th century. Instead of despairing, one is thus invited to rejoice with the demise of reductionism and celebrate the pandemonium of possibilities that is bound to sweep all of physics. Philosophers, but even physicists (e.g. Ellis 2012) have entertained the possibility that reductionism is both inaccurate as a descriptive device but also undesirable on methodological grounds. For example, Cartwright 1999 has

³⁹Craig & Koren discuss how an EFTist would describe the universe with this complication and claim that the IR divergences could be absorbed by the introduction of a new degree of freedom. Nonetheless, it is not clear that this new field should be taken seriously as representing something physical or as a mere computational artifact, much like ghosts in non-abelian gauge theories.

explicitly endorsed a disunified picture of science where different levels of reality and types of phenomena do not reduce to one another within a (naively conceived) Maslow-like pyramidal structure. Apart from its inaccuracy, Cartwright finds that the more fundamentalist nature of theorising implied by the standard reductionist picture actually hurts progress in areas such as economics and is thus to be rejected, if some of the pathologies in that domain are to be cured. Similarly, we already saw Cao and Schweber 1993 argue for some form of radical anti-foundationalism in light of EFTs, while a more top-down conception of science was also painted with broad brushes by Adlam 2019).

While all these ideas are interesting and intriguing for their own sake, we will not examine them here as possible lessons to be drawn from either Wallace's project or the EFT framework in general. The reason is two-fold. On the one hand, I think that there are plenty of reasonable alternatives that we can pursue without digressing so much from our engagement with EFTs. On the other hand, I think that escaping Wallace's dilemma through such a radical departure from our prevalent understanding of 20th century physics will simply play into his hands. A more conservative or serene approach is, as I hope to have shown, more than adequate to suppress the problem.

Lessons for Reductionists? The above considerations notwithstanding, there are a few things reductionists need to take into account in light of the restrictiveness of the EFT framework and/or (unnatural) decoupling. For example, if we follow Weinberg 1987 in his trichotomy of reductionism into:

1. Theory Reductionism: deny autonomy to special fields - all to be absorbed into fundamental (high energy) physics
2. Explanatory Reductionism: progress in higher-level theories requires progress in lower-level theories
3. Objective Reductionism: convergence of "arrows of explanation" to fundamental physics

we see that breakdowns of the EFTs can be seen as entailing a failure of the third "objective form" of reductionism as well⁴⁰. The problem, of course, is that the delicate connection between high-energy and low-energy might seem as a threat to the neat distinction between more

⁴⁰Weinberg himself grants the arguments made by "opponents" of reductionism against the first two, unnecessarily strong, forms of the view. Accordingly, we will not comment on these any further here.

fundamental and less fundamental ontological levels - in effect belying the assumption that acted as the thrust of progress in physics throughout the 20th century. Yet, a complete departure from the merits of the reductionist agenda is in all likelihood too hasty a move. What is warranted, instead, is an appropriate evaluation of some of its central tenets:

- Ontologically: reductionism is very much alive in the sense of compositionality: particle physics still reveals to us the building blocks of reality (perhaps up to and within some given scale(s)) so that the whole-part relationship in a compositional sense is preserved
- Theoretically: the EFT framework still applies as we know (under appropriate conditions); even if naturalness is violated, some form of decoupling is ensured whenever we can “summarise” the physics of a lower level in the parameters of an effective theory; we can still track the levels of reality within the (manually) constructed tower of EFTs
- Fundamentalism: this aspect of reductionism, construed as the possibility of finding a unique underlying theory from which to “derive” the rest of physics from the bottom-up, does not find its proper home within the EFT framework; if violations of naturalness are indicative of some form of UV/IR mixing, then the ontological picture emerging for any successor theory will probably subvert our expectations⁴¹
- Predictability: this might also need to be restricted in the case of unnatural EFTs; one should be content with “cohesive” as opposed to bridging principles between the various theories; some “fundamental” parameters might simply be adjusted to specific values for no reason other than as the input consistent with our empirical data

All in all, reductionism, much like realism, should be treated in a more methodological manner as representing a good strategy for understanding reality within some specific regimes. Unnatural EFTs, mixing of UV/IR etc do not signal its complete demise but the loss of its universal appeal as *the* methodological doctrine on which to base our treatment of all phenomena in nature. Some times the optimal strategy is to stop digging.

⁴¹Again, this does not force us to go all the way with Cao & Schweber’s more “radical” or “enthusiastic” appraisal of the program. See chapter II2b.

3.6 Conclusions

To wrap up this rather exhaust...ive chapter, let's highlight the main results of our investigation. We saw that the hierarchy problem is an instance of a failure of naturalness. We identified the prevalent conception of this principle in the literature in the form of technical naturalness which dictates that naturalness violating terms can be accepted if their elimination leads to an enhanced symmetry of the theory. We identified three distinct stances one can take with respect to the problem: a) acceptance, b) denial and c) revision. We first examined some attempts at resolving the problem with the introduction of new physics. We saw that most of the proposals have failed to be experimentally confirmed or lie on shaky conceptual ground. This led us to skepticism about the usefulness of naturalness as a guiding principle. This suspicion was reinforced when we found out that most examples of successful application for naturalness were reconstructions of previous discoveries as opposed to original predictions made on its basis. In cases in which naturalness was associated with some epistemic risk, it has failed to lead to breakthroughs. If we view it through the methodological lens of realism we presented in the previous chapter, naturalness fails to pass the test. So far it might at best be treated as a tentative heuristic for deliberating on speculative extensions of physics.

It would be premature to end our investigation here, however. To refute a proposed principle it is necessary to find fault with its conceptual foundation. Quite naturally, then, our next stop was to examine dissenting voices arguing against the predominant trend of taking naturalness seriously. We found that formulations of naturalness as aversion to fine-tuning were mostly circular as they heavily depended on our intuitions about what constitutes a typical measure on the space of bare parameters. Worse still, we discovered that the physical significance of unnatural parameters quickly evaporates under appropriate regularisation and/or renormalisation schemes. These concerns gave rise to an investigation over the motivation behind naturalness, which eventually revealed a potential connection with the autonomy of scales requirement. Irrespective of whether the majority of the proponents of naturalness have explicitly construed the principle in this manner, we realised that this is its strongest formulation. This is because it directly connects it to EFTs – with violations of naturalness being seen as failures of the EFT framework.

Now, it is pretty clear that repercussions of naturalness violations be pretty damning if one grants that $D = N$, i.e. naturalness = decoupling. In fact, if more extreme lines are followed the problem might even beget a revision of the methodological approach to modern physics!

Wallace 2019 is the most vocal expression of the problem going as far as to equate violations of naturalness with failure of reductionism. After analysing his argument, we came to understand that naturalness can be understood along two axes: one reflecting its scope (i.e. whether it permeates the whole of physics or just EFTs) and its relation to decoupling (i.e. whether it is to be seen as equivalent to it). We can summarise the dilemma we face as in 3.3:

Table 3.3: The Decoupler’s Game

NATURALNESS is...	DECOUPLING	NOT DECOUPLING
UNRESTRICTED	Wallace	Bain (?)
RESTRICTED	Physicists Williams (?)	“Dissenters” (Rosaler, Harlander Hossenfelder)

The road to resisting Wallace’s argument and mainting violations of naturalness without compromising the effective reading of QFTs is now clear. On the one hand, one can deny naturalness the status of a universal or unrestricted principle. Arguing against taking the analogy between statistical mechanics and quantum field theory too seriously, one can undermine the importance of cut-off sensitivity (as in the Wilsonian approach to renormalisation) of parameters like the Higgs mass. While the techniques employed in both frames are similar mathematically, they do not necessarily reflect a deeper physical congruity between the two fields. The continuum version of the renormalisation group, which treats the cut-off dependence of parameters as purely conventional (and is the technique actually employed by physicists), could be seen as further evidence for this. On the other hand, we have also seen that the naturalness and decoupling should not be seen as equivalent. A more appropriate relationship is one of genus-species with natural decoupling being only a special form of decoupling that obeys additional constraints on the values of parameters. More encompassing notions of decoupling such as those presented by Bain [“heuristic decoupling”] or Franklin can account for the applicability of EFTs without imposing the (additional) naturalness restriction. That these approaches are on the right track is empirically corroborated: so far it has been possible to extract perfectly sensible results from the standard model of physics in spite of the unnatural Higgs mass parameter. In this sense, reductionism (perhaps a less ambitious form of it) is perfectly safe.

Finally, we can, and indeed did, contemplate the possibility of a failure of decoupling (equivalent to naturalness) and saw that far from signaling a problem for all of physics (as suggested by Wallace), it might be construed more “simply” as an indicator of the restrictiveness of the EFT framework. It is often mentioned in the physics literature that every theory can be treated as an effective theory. This, however, conceals several nuances of the word “effective”. If “effective” means that theories only apply to certain phenomena and might produce nonsensical results when applied outside their proper regimes, then any theory is indeed effective. If “effective” also implies that standard renormalisation group techniques (with a clear-cut separation between UV and IR degrees of freedom) apply, then some theories might not be accurately treated as such. Indeed, having briefly examined non-commutative field theory we saw that UV/IR mixing, which in no way is an outlandish scenario, violates the expectations (standard assumptions) of the EFT framework. As we shall see in the next chapter, further evidence of the limitations of EFTs might come from cosmology and QFT in curved spacetime. Although, this might not directly lead to new physics (that is, new fields), it might lead to a revised stance to the way we understand physical theories in certain regimes. Wallace in this case would be right insofar as we understand his claim to be that new techniques would need to be invented to deal with what lies in the “great beyond”. He would still be wrong about extrapolating the problem to all of physics.

What of the hierarchy problem itself, then? This is a trickier question to address, but it is worth connecting it to our discussion of methodological realism in the previous chapter. There, we saw that the various theoretical elements such as entities, principles, mathematical techniques posited and employed by a theory are ultimately vindicated, i.e. are accorded a realist status, when they let us enhance the inferential capacity of our theories and lead to novel expectations about the “inner workings” of nature. If we adopt this mindset, what we can say of the hierarchy problem is that insofar as it relies on a naturalness assumption it should not be accorded the significance it has been given so far. This does not mean that the problem has no heuristic value in motivating the construction of models and the development of new physics (witness the wealth of results in extensions of the standard e.g. through SUSY), but one must always remain vigilant about the shaky physical grounds on which it lies. The elevation of a heuristic to a physical principle should involve instances of unambiguous successful application in guiding progress in science. Mere analogical arguments cannot be a substitute for that.

Table 3.4: Solving the Decoupler’s Game

NATURALNESS is...	DECOUPLING	NOT DECOUPLING
UNRESTRICTED	reject QFT – CSM analogy	assess empirically
RESTRICTED	reject $N = D$ equation	\emptyset

This table presents the possible ways naturalness can be rejected without compromising the methodological appeal of reductionism or the validity of the EFT framework. If restricting the problem to QFT is not accepted, one can reject the equivalence of naturalness and decoupling. Decoupling, in the form one comes across in QFT, can also be considered as a more special case of reductionism and need to challenge its usefulness in the rest of physics.

Chapter 4

When EFTs Seem to Fail II: Cosmology

The hierarchy problem is not the only indication that there is something unsatisfactory about the SM or perhaps the EFT framework itself. Physicists have identified yet another anomaly in our attempts to probe physics beyond the SM – this time in the critical junction where QFT meets GR. This is the (infamous) cosmological constant problem, which, according to conventional wisdom, most vividly shows the predictive inadequacies of QFT. Interestingly, however, given its dual origin, the problem can be seen as either a problem for GR or SM – or even both. It comes as no surprise, then, that some attempt to cure the problem by modifying gravity, others by including new fields in QFT or by transitioning to a quantum theory of gravity.

Our main focus here is in the CCP's impact on the EFT framework and thus we will be mostly preoccupied with the aspects of the problem that have something to reveal about the merging of GR and QFT. In this analysis, we will pick up the thread we left in the previous chapter about possible limitations of the EFT framework. We will find that the challenges facing QFT on curved spacetimes highlight specific assumptions of this framework that may not be extended to more generalised spacetimes (i.e. of a non-Minkowski kind). With these assumptions undermined, the very preconditions necessary for the applicability of the EFT framework are undermined as well – a clear signal of its potential breakdown. An example of this, as we shall soon enough see, is a possible failure to distinguish between negative and positive energies or unambiguously define a vacuum state.

The plan of this chapter is the following. In the first section we introduce those aspects of QFT and GR that are relevant to the CCP. In particular, we discuss the role of the vacuum in standard formulations of QFT, examine phenomena that (allegedly) point to the existence of some form of vacuum energy as well as the role the cosmological constant term plays in

cosmology. In the following two sections we briefly describe how the CCP emerges breaking it down with the use of a decision tree. We offer a rather terse assessment of the problem as standardly presented referring to the literature for more developed versions of the criticisms. In the final, but most important, section we return to EFTs and re-frame the CCP as analogous to the third formulation of the hierarchy problem. We find that, while strong analogies between the two cases hold, the CCP offers the best ground for appreciating the shortcomings of the EFT framework and, by extension, more urgently calls for a careful examination of the conditions lying at its foundation.

4.1 The Void, the Vacuum and Nothingness

Most evident in philosophy's infancy years and the work of Parmenides, fascination with nothingness never really waned: this idea made periodic reappearances in the guise of a matter-less vacuum, the spatiotemporal void or the state of non-existence par excellence. Philosophers as different as the ancient atomists, Aristotle, early modern dualists and materialists (think of Descartes, Boyle, Leibniz, Gassendi et al) or even existentialists like Sartre(!) have written extensively on the topic. Yet, for all the intellectual power invested in uncovering the mysteries of the void, the concept to this day remains largely elusive. In a rather stunning reversal of fortune, it was modern physics that vindicated, to some extent, the philosophers' earlier aversion to the vacuum: QFT, which treats the vacuum as a state of minimum energy, shows that what was conceived of as emptiness might be a state of existence more active than previously imagined. The purpose of this section is to introduce the two main players in the CCP drama: the QFT vacuum and the cosmological constant term in GR.

4.1.1 The Abhorred Vacuum

Early attempts to tackle the nature of what we might today call the vacuum stumbled upon the paradoxical aspects of non-being. Broadly speaking, Parmenides' work on the distinction between being and non-being left subsequent natural philosophers with a dilemma: either reject the idea of the vacuum tout court or accord it some peculiar form of existence. As is well-known, the former stance dominated the intellectual scene from the work of Aristotle to that of early modern philosophers with its abhorrence to the vacuum. Descartes and Leibniz famously argued against its intelligibility advocating for an extended plenum in its stead. Dissent has

always been present, however. The atomistic tradition stemming from the work of Democritus, Epicurus treated emptiness as a state of non-being in which individual loci of existence, the atoms, would roam and combine to give rise to the world we see around us. This flame was kept alive throughout the ages and would finally exact its vengeance after millennia of subjugation.

Modern physics is, with good reason, understood to have vindicated the atomistic tradition by affirming its fundamental insight about the existence of atoms and a state of non-being roughly corresponding to unoccupied space. Experimentalists such as Torricelli and Boyle claimed to have been able to extract all air from tubes using vacuum pumps - applying pressure to the idea that nature would prevent such states from coming to being¹. Further progress in the theory of electromagnetism, with the introduction of force fields, seemed to annul this result, however. Instead of pure nothingness, empty space would now be pervaded by fields of electromagnetic, gravitational or any other (unknown) force fields affecting any wandering matter².

Quantum Mechanics As expected, quantum mechanics and quantum field theory further complicated things. On the one hand, simple QM systems, such as that of a particle trapped in a potential well, most clearly illustrate the existence of (ground) states of minimal energy. This is the case for the most celebrated of all QM systems, the simple harmonic oscillator (sho), which is described by the Hamiltonian:

$$\hat{H} = \frac{\hat{p}^2}{2m} + \frac{1}{2}m\omega^2\hat{x} \quad (4.1)$$

The ground state for the sho, as is well-known, is written in terms of Hermite functions and its energy is:

$$E_0 = \frac{1}{2}\hbar\omega \neq 0 \quad (4.2)$$

When everything is recast in terms of \hat{a} and \hat{a}^\dagger , the excited states can be obtained by acting on the ground state with creation and annihilation operators. So, starting from a state of minimal energy, one can describe the evolution of the sho system as transitioning from one energy level to the other.

¹The atomistic hypothesis itself would re-enter the scientific realm with Dalton's work in chemistry, although its establishment as a mainstream view would have to wait till Perrin's work in early 20th century. All of this will not concern us here, though.

²Interestingly enough, 19th century physics witnessed the sprout of programs intended to account for matter in terms of the ethereal medium which was assumed to be the carrier of the electromagnetic field. For a very interesting compilation on these abandoned projects see Kragh 2011.

The standard story one is then told is the following. Even though only differences between energy scales are observable, the sho at its ground state possesses some minimal, not necessarily zero, energy. This is typically interpreted as a consequence of the uncertainty relations: the minimal energy represents the “restless” disposition of the system to defy attempts at specifying its position and momentum at the same time. Whether and how this story makes sense and whether and how one wants to adopt is probably a controversial issue, but accepting that it somehow proves the reality of vacuum energy is probably premature. The reason is that the ground energy does not possess an absolute value, but can easily be redefined by adding a constant c :

$$E'_0 \longrightarrow E_0 + c \quad (4.3)$$

As a result, all energy differences can be redefined in order to accommodate a zero value for E_0 . The essential condition that needs to hold is that $E_{QM} - E_{CM} > 0$, i.e. the classical and quantum energies should differ, without specifying a precise value for either. So, there is some freedom in specifying the exact values of the parameters.

Quantum Field Theory Is the above reparametrisation a sensible move to make? Well, as it happens, this is precisely what one does in QFT when dealing with the first and most rudimentary form of infinity one comes across. When we write down the Hamiltonian for a free scalar theory and we use the commutation relations for a, a^\dagger :

$$\hat{H} = \int \frac{dp^3}{(2\pi)^3} \frac{\omega_p}{2} (\hat{a}_p \hat{a}_p^\dagger + \hat{a}_p^\dagger \hat{a}_p) = \int \frac{dp^3}{(2\pi)^3} \frac{\omega_p}{2} (\hat{a}_p^\dagger \hat{a}_p + \frac{1}{2} \delta^3(p)) \quad (4.4)$$

we notice that the last term leads to an infinity once the integral over momenta is performed. However, this can easily be eliminated if we apply normal ordering, i.e. if agree to set all creation to the left of annihilation operators. In this manner, we simply(!) shift the energy by an infinite constant equal to the above divergent integral to obtain a finite result consistent with our experience³.

In fact, the difference between the two cases (QM and QFT) is only quantitative. After all, quantum fields, when viewed through the QFT lens, can be described as huge configurations of superposed states of quantum harmonic oscillators at different energy-momenta. Instead of dealing with a simple discrete system with a finite number of degrees of freedom, we are here

³One might remark, with good reason, that this is what we do with a theory's parameters when renormalise it.

considering a continuous system with an infinite number of degrees of freedom - whence the infinite divergence. The process of regularisation through normal ordering neatly separates the infinite from the finite part in a strategy closely analogous to the shifting of the sho ground state by a constant.

In both cases, the arbitrariness of the procedure is attributed to the fact that claims about absolute values of energy –as opposed to differences in energy levels– lack physical meaning. One can adjust these parameters anyway they want as long as these lead to correct estimates for the physically meaningful quantities. The situation changes when gravity enters the picture, however. As Einstein’s theory has taught us, any form of energy gravitates. Should this include the infinite energy we just regularised away? There is more to the story, it seems.

4.1.2 The Infamous Constant

So much ado about nothing in QFT. Nothingness does not just have to do with the state of limbo characterising field excitations at the microscopic level. If we turn to GR, we can also extract interesting facts about nothingness as the “void” or empty spacetime. The classic story here has to do with Einstein’s infamous “greatest blunder” or the Λ (or cosmological constant) term in his field equations [EFE]⁴:

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} \quad (4.5)$$

Particle physics has no monopoly on standard models: cosmology also flaunts its very own Λ CDM model! According to this, the universe is essentially a multi-flavoured cake consisting of matter and (cold) dark matter, radiation and whatever is represented by the Λ term – given the captivating name of “dark energy”. The evolution of the universe goes through phases that are determined by the density of each ingredient as compared to a critical energy density corresponding to a flat universe. The Λ term represents a peculiar sort of fluid behaving in

⁴There is a lot of confusion over Einstein’s motivation for the introduction of the cosmological constant and the significance he attributed to it. In most popular accounts, Einstein simply inserted the Λ term by hand in order to ensure that a static universe would come out as a solution from his theory - the blunder of course being that he failed to use the instability of his solution to predict that the universe actually involves. This attributed (by Gamow) quote contains a grain of truth about Λ but misconstrues the original problematic that led to it. Einstein was not so much interested in eliminating the possibility of an evolving universe. After all, why would the scientist behind the most spectacular revolution of our very understanding of space and time be intimidated by the possibility of an evolving universe? His primary interest was to render GR adhere to what he thought of as Mach’s principle, very roughly the idea that inertia can be accounted for in terms of the matter content of the universe without invoking some background structure.

an anti-gravitational, i.e. repulsive manner. According to established wisdom, the universe is currently dominated by this fluid which also drives its expansion. The physical meaning of this term remains elusive, however. Its interpretation directly influences the nature and importance of the CCP itself.

4.2 What is the Problem with the Cosmological Constant?

Anyone ignorant of basic chemistry would arguably be surprised to learn that a benign substance such as salt is created by some of the most “violent” elements. The CCP can be seen as the inverse kind of surprise: two rather innocuous elements combine to produce a most disturbing result. As long as we disregard the influence of gravity, bubble diagrams play no role and any infinities associated with the vacuum simply drop off our calculations. Similarly, as long as we forget about the zero-point energy, GR can pretty accurately describe the evolution of the universe. As soon as we attempt to combine the two, however, we quickly realise that peaceful co-existence is not an option. Here is the standard story why.

Einstein’s field equations (EFE) describe the way matter, represented by the stress-energy tensor, $T_{\mu\nu}$, influences spacetime geometry. Since $T_{\mu\nu}$ can be defined using the action of the theory of our interest⁵, we may start by adding, for the sake of simplicity, a scalar field into the action for the gravitational field:

$$S = \frac{1}{2\kappa} \int d^4x \sqrt{-g} (R - 2\Lambda_B) + S_{matter}[g_{\mu\nu}] \quad (4.6)$$

As a result, the EFE will eventually see the addition of a vacuum term

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + \Lambda g_{\mu\nu} = \kappa T_{\mu\nu} + \kappa \langle T_{\mu\nu} \rangle \quad (4.7)$$

We are now a relabelling away from seeing the problem. By redefining the cosmological constant as:

$$\Lambda' \longrightarrow \Lambda + \kappa \langle T_{\mu\nu} \rangle \quad (4.8)$$

⁵It is important to appreciate the ambiguities in this step. Even in the context of classical field theory, there are various procedures to come up with $T^{\mu\nu}$ which in the unproblematic context of electromagnetism lead to the same result. Such procedures might involve defining the stress-energy tensor in ways that are not equivalent in general. Indeed, as it turns out, standard procedures such as those by Hilbert and Noether, which give equivalent results in the context of EM, lead to inequivalent results in the context of spin-2 theories. See Baker, Kiriushcheva, and Kuzmin 2021 for a technical discussion of these complications in the context of higher-order linearised (Gauss-Bonnet) gravity.

we notice that the cosmological term receives contributions from two distinct sources: one from the original, bare or built-in term Λ and one from the fluctuations of the new matter field⁶. This is where trouble arises. As soon as we write the EFE like:

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + (\Lambda_B + \Lambda_Q)g_{\mu\nu} = \kappa T_{\mu\nu} \quad (4.9)$$

we can't help but wonder about the possible values of the $\Lambda_B + \Lambda_Q$ term which, apparently, receives contributions from the bare value of the constant and the fluctuations of the vacuum.

The Argument It is instructive to lay out the argument in a detailed, formal manner to more clearly see the various assumptions involved:

- P1 In GR, the EFE can be supplemented with an additional term containing an unspecified parameter, Λ .
- P2 When solving the equations to obtain evolutionary models of the universe, we notice that this additional term should be responsible for some kind of expansion.
- P3 The universe is indeed expanding. The parameter is set consistent with observations equal to $\Lambda = 1.1056 \times 10^{-52} m^{-2}$.
- P4 We must attribute some cause (call it “dark energy”) to this expansion and then the question naturally arises: whence this repulsive force?
- P5 In QFT the various fields are now contributing to the EFE influencing the geometry of spacetime. All known SM ingredients behave attractively. Alternatives include:
 - R1 Add an additional (quantum) field that behaves repulsively and whose expectation value is such that $\langle 0| T_{\mu\nu} |0\rangle = \Lambda$.
 - R2 Connect the new anti-gravitating form of energy with the ground state energy of quantum fields.
- P6 redQFT contributions are of the form $\langle 0| T_{\mu\nu} |0\rangle = -\rho_{vac}g_{\mu\nu}$ and, at the lowest estimate, of the order 10^{10} or even as high as 10^{70} .

⁶In this brief presentation of the argument we are mainly following Martin 2012 in his (much more) detailed discussion. Martin is making a distinction between a classical and quantum version of the CCP. The quantum version, which involves contributions from the ground state of quantum fields, is the focus of our attention as we are interested in the “marriage” of QFT and GR. The classical version of the problem is about the changing value of the classical field potential -corresponding to the classical Λ - between phase transitions.

C1 To obtain a value consistent with empirical data a very delicate cancellation (fine-tuning) between the QFT and bare term is required.

C2 Therefore, unless we are prepared to admit fine-tuning, some modification of our current physics is needed.

Assessment We will critically examine the assumptions of the argument. To systematise our analysis we will be using the following decision tree:

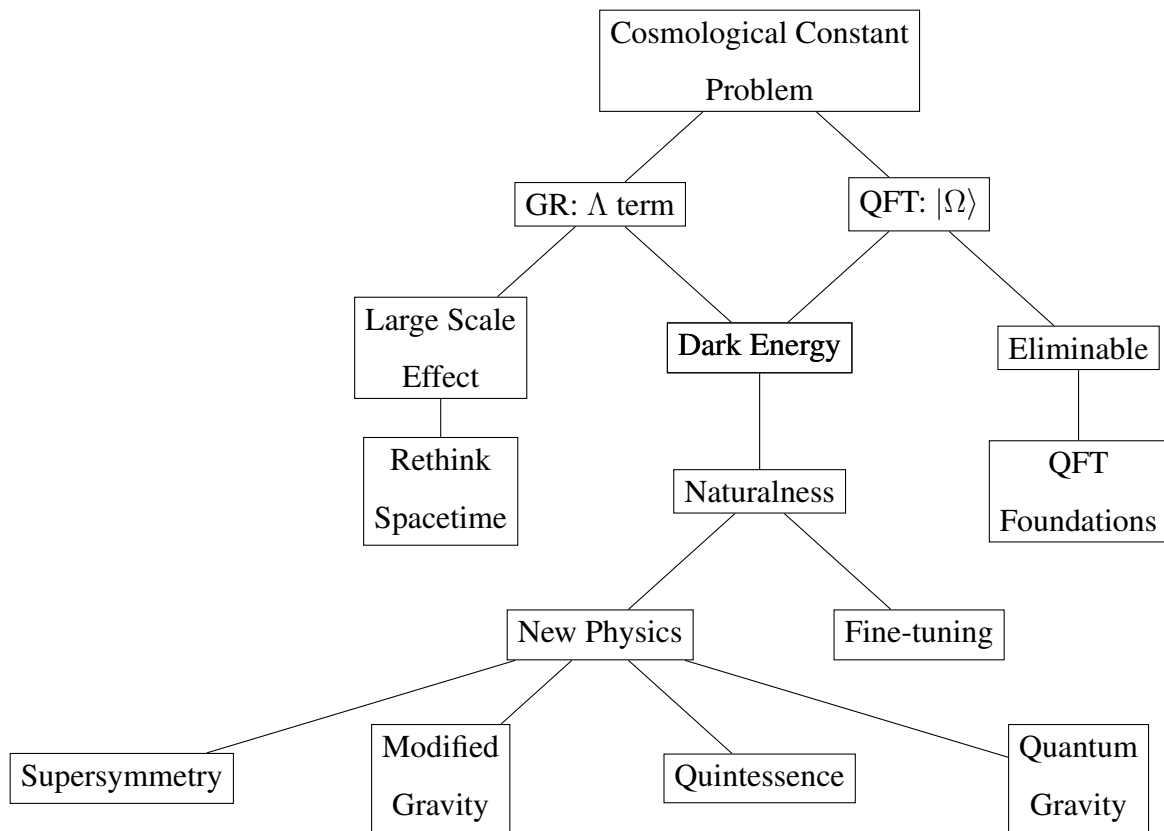


Figure 4.1: The Cosmological Constant Problem: A Decision Tree

4.2.1 Elastic Spacetime

Perhaps the most straightforward way to dissolve the CCP is by treating the expansion of the universe as a large distance effect of the gravitational interaction. Essentially, the move here is to remove the shroud of mystery around dark energy by taking the cosmological term as an inherent feature of GR. As Einstein once remarked the two sides of the EFE are made of different materials: the LHS out of fine marble of geometry while the RHS out of the dull wood of matter. There is an inherent ambiguity as to the place of the $\Lambda g_{\mu\nu}$ term in EFE. The received

view is to treat as part of the matter content of the theory (hence the connection with QFT). On the alternative examined here, however, the term must be treated as part of the geometry, i.e. a feature of spacetime.

How should we think of this? Perhaps the best way is to concede that gravitational interaction is quite unlike what we have conceived of so far: while primarily seen as attractive, this has been due to the fact that our intuition was formed, at least up until the 20th century, by small distance phenomena. Gravity, we thought, is the one-sided interaction par excellence. Instead, gravity is more like a two-faced Janus: one pair of eyes set on attraction at small enough distances and another set on repulsion at scales large enough to be comparable to cosmological scales⁷. Alternatively put, when we say that the universe is expanding, we are not talking about a mysterious form of energy that stretches space and time, but rather about an intrinsic property of spacetime itself: spacetime turns into an “elastic” entity⁸!

Now, the move might at first glance seem annoyingly ad hoc in that we endow gravity with an additional feature just to make sense of an observed phenomenon. However, there are at least two reasons to shake this feeling. One reason is to recall the fact that GR brought about a dramatic reconceptualisation of gravity that defied both our everyday intuitions and the established physical framework of its time: after GR, gravity would no longer act at a distance, spacetime itself would be curved, forces would be geometrised and so on. In this vein, perhaps another unknown fact about gravity is this large distance effect undetectable at distances below cosmological scales. To lend further credence to this point, consider now a stronger point: other than historical contingency, there is no reason against including a cosmological term in EFE. Einstein’s “blunder” was that by adding this additional unspecified parameter he would be able to predict an expansionary or contractionary universe without changing the physical content of his theory. In fact, one cannot help but wonder: had the constant been introduced during the first days of GR and subsequently proven to lead to a clear prediction on behalf of the theory, would we be still talking about a fundamental problem with the cosmological constant today?

⁷An interesting, albeit loose, historical analogue to this alternating nature of the (gravitational) interaction is the theory of matter as force by Boscovich and Child 2018, who took the fundamental force governing the universe to exhibit a similar dualistic (repulsive and attractive) behaviour.

⁸It is not clear that adopting this view needs to commit someone to a substantialist or absolutist position on the ontological status of spacetime. The relationalist will probably want to argue that this “repulsive” push is built-in the gravitational field itself and no additional spatiotemporal substance needs to be assumed – much like the EM field and ether. We will assume that one can remain neutral about this debate and simply not pursue it here.

4.2.2 When the Nothing Nothings

If only life were that easy, however! Even if the above allows us to lend credence what Bianchi and Rovelli 2010 suggest in their first conclusion (for the restricted context of GR):

...the cosmological constant term is a completely natural part of the Einstein equations. Einstein probably considered it well before thinking about cosmology. His “blunder” was not to add such a term to the equations: his blunder was to fail to see that the equations, with or without this term, predict expansion. The term was never seen as unreasonable, or ugly, or a blunder, by the general relativity research community. It received little attention only because the real value of λ is small and its effect was not observed until (as it appears) recently.

we will still face a problem when QFT enters the scene. This is because in a semi-classical approach, quantum fields will contribute to the average value of the stress-energy tensor:

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + \Lambda g_{\mu\nu} = \kappa \langle T_{\mu\nu} \rangle \quad (4.10)$$

Of course, by itself this is no problem. Setting aside qualms about the coherence of the semi-classical approach to gravity for the moment, the quantum fields will determine the geometry of spacetime quite similarly to ordinary matter one inputs in $T_{\mu\nu}$. Unfortunately, ordinary matter will only contribute to the contraction of the universe. Something further will be needed for the “outward” repulsion. What can this “something” be? a) One option is to introduce a new, yet to be discovered, field with the right equation of state, i.e.

$$p = -\rho \quad (4.11)$$

where p is the pressure of the “dark fluid” and ρ its energy density. This new field would fulfill the role of “pushing” spacetime outwards against the contractive effect of radiation and matter. b) Another option is to consider the effect of zero point energy, i.e. the energy “latent” at the ground state of the contributing quantum fields, as the “something” cause the expansion. It is this second option that results in cracks to the nice picture we painted above.

The reason is simple. If, as is widely accepted in physics today, vacuum energy exists, then, as any other form of energy, it must be gravitating. Of course, there is no problem when we do not take gravity into account, because when calculating probability amplitudes using Feynman

rules, we simply drop bubble (or vacuum diagrams) without penalty. This was the very first step in the simplification procedure of the path integral:

$$Z[\phi] \longrightarrow W[\phi] \longrightarrow \Gamma[\phi] \quad (4.12)$$

Now, however, we cannot simply repeat this trick and disregard these diagrams because all matter couples to the gravitational field. The same should apply to vacuum fluctuations. At the same time, we know that the actual value obtained from all these diagrams is infinite! How can we extract any meaningful estimate of the cosmological constant? The standard, rather clumsy way, of obtaining a preliminary finite answer is by using cut-off regularisation with a smooth cut-off⁹ set at M so that:

$$\Lambda \sim \frac{1}{4\pi^2} 6M^4 \quad (4.13)$$

This is a quadratic dependence on the value of M , which will certainly make the value of Λ explode even when a pretty small cut-off is used. If we are as bold as to take $M = M_{Pl} \sim 10^{19}$ GeV the value of the cosmological constant becomes $\Lambda \sim 10^{76}$ GeV or about 120 orders larger than what is measured - see (P3). Even if we do not take the cut-off to be as high as the Planck scale, but rather closer to the electroweak scale $M_{EW} \sim 10^2 - 10$ GeV, (the idea, of course, being that new physics might be around the corner at energies far below the enormous energy levels of the Planck scale), the cosmological constant is still off by around 50 – 60 orders of magnitude. A disaster either way!

Lest this be an artifact of the regularisation scheme, one will be well-advised to perform the calculation using an alternative scheme, such as dimensional regularisation. In fact, as Martin 2012 (p. 12) notes, there is a deeper ongoing issue here. As we know, cut-off schemes usually violate certain symmetries of a system like Lorentz invariance. In the case we examine, calculations lead to an equation of state:

$$\langle p \rangle = \frac{1}{3} \langle \rho \rangle \quad (4.14)$$

which is the equation of state for radiation as opposed to that of a “repulsive fluid”! This turns out to be an artifact of the regularisation scheme since dimensional regularisation gives a

⁹See section IV in Martin 2012 for detailed calculations of the mean energy density $\langle \rho \rangle$ and the mean pressure $\langle p \rangle$ for the stress-energy tensor of a scalar field Φ . Martin performs the calculations using both a hard and a smooth cut-off (before turning to dimensional regularisation). We only quote the latter result here as the main feature, i.e. the quartic dependence on the cut-off scale is present in both.

result consistent with the right equation of state $\langle p \rangle = -\langle \rho \rangle$. Nevertheless, despite the fact that estimates for the energy density under this scheme differ:

$$\Lambda = \frac{m^4}{64\pi^2} \ln \left(\frac{m^2}{\mu^2} \right) \quad (4.15)$$

with μ an introduced energy scale, the problem is not yet resolved as the value is off by several orders of magnitude. Regularising the problem away is not a straightforward business. Koberinski 2021 has recently tried to impose conditions on the kind of regularised values that we are warranted to take seriously. He argues that a key requirement for any quantity to count as a prediction in QFT is that it satisfies what he calls “Minimal Requirements for Candidate Predictions”:

- (1) the quantity is largely insensitive to the regularization procedure; and (2) it is largely insensitive to changes to the value of the regulator.

Since the regularised cosmological constant estimate violates these conditions, Koberinski argues that the values obtained should not be trusted. I think that the point about the CC falling outside the scope of QFT is apt, but one should also be sure not to overlook the most important aspect of the CCP as problem – a point to which I turn in the end of this section.

If the above numerical estimates for the value of vacuum energy are approximately correct, then as soon as quantum fields are turned on, their vacuum energy contributions will make it very hard, i.e. it will require a unacceptable level of fine-tuning, to make EFE compatible with observational data¹⁰. Indeed, assuming that the vacuum energy of quantum fields is gravitating, the bare or classical part of the Λ term need to be set with extreme precision to cancel the quantum part in a way that leaves us with the right empirical value:

$$\Lambda_B \sim \Lambda_{EXP} - \Lambda_{QM} \sim 10^{-52} + 10^{76} = 10^{76}(1 + 10^{-128}) \quad (4.16)$$

- an adjustment at the 10^{128} decimal will be needed. Obviously, things can get a bit better if the estimate for Λ_{QM} is lower, but fine-tuning would still be very high even at scales as low as 10^3 GeV. One can easily appreciate the analogy with the Higgs boson case we examined in the previous chapter. The important difference in this case, however, is that the CCP is not an

¹⁰Problems are further exacerbated if the cosmological “constant” turns out to be a non-constant, i.e. evolving in time. In this case, one would need to impose different values on the bare parameter between phase transitions just to accommodate the relevant data.

internal problem of a specific theory but rather the offspring of the attempted combination of two distinct frameworks. This exposes the CCP to further criticism perhaps unavailable in the case of the hierarchy problem.

4.2.3 Turning on the Vacuum... Cleaner

So far we have lent support to Heidegger’s idea that “the nothing nothings”. But, surely, that “nothingness” might somehow “exist”, let alone “do” anything seems a highly absurd suggestion. Since premises (P1) – (P3) cannot be really disputed and the possibility of an “elastic” spacetime specifically targeting (P4), has been examined, we now turn to premises (P5) and (P6), which involve the QFT part of the CCP. Skeptical arguments in this direction challenge the traditional accounts for the reality of vacuum energy and/or the well-definiteness of the QFT conceptual framework on a non-Minkowski background.

Questioning the Received Wisdom Even if one is content with a picture in which Λ appears as an additional parameter in GR and deems aspirations to uncover its origin as at best unnecessary, the reality of the QFT vacuum energy will, as we just saw, eventually come back to haunt them. It does not matter whether one believes in equating Λ in GR with whatever vacuum energy contribution comes from QFT. As soon as this QFT contribution makes its appearance in the EFE, the CCP follows.

Assumptions: Nonetheless, the assumption that the vacuum energy is a real gravitating form of energy contributing to the EFE is far from innocent as it makes certain assumptions:

1. **Vacuum energy:** The vacuum is not a state of non-existence: it is replete with quantum fluctuations, represented by bubble diagrams that carry ontological significance and actually influence physical interactions.
→ justification: Casimir effect, Lamb shift effect
2. **Equivalence principle:** Not only does vacuum energy exist, but it also gravitates!.
→ justification: strong equivalence principle, all forms of energy must interact with matter - and vacuum energy should be no different.
3. **Stress-energy tensor form:** for matter to contribute a cosmological-constant-like term in EFE the stress-energy tensor needs to be of the form $T \sim g_{\mu\nu}\rho$; essentially (P6)

→ *justification: QFT framework, equivalence principle*

4. **QFT in curved spacetime:** underlying the whole discussion of course is that the present framework of QFT makes sense or is (approximately) accurate in some non-Minkowski background

→ *justification: TINA, effective description*

Assessment: There are reasons to reject or at the very least be skeptical of each of the above assumptions¹¹. First, despite frequent claims in favour of the reality of vacuum energy, its existence is not an established or uncontroversial fact. Typically, evidence in favour of the vacuum in QFT comes mainly from phenomena like:

- *Casimir Effect:* the attractive force arising due to vacuum fluctuations between two metallic plates placed at very close proximity
- *Lamb Shift:* the energy discrepancy between the $^2S_{1/2}$ and $^2P_{1/2}$ states of the hydrogen atom attributed to interactions of the electron with the fluctuating vacuum
- *Bubble Diagrams:* these diagrams, mathematical descriptions of processes involving “virtual” particles, are physically interpreted as fluctuations of the vacuum

An important blow against the orthodox explanation of the Casimir effect was dealt by Jaffe 2005, who presented an alternative account for the attraction between the plates in terms of van der Waals forces¹². A plethora of papers sceptical of the reality of the QFT vacuum have also assailed the cogency of the “Casimir effect → reality of vacuum energy” derivation. For example, apart from the paper by Jaffe, one can consult Rugh, Zinkernagel, and Cao 1999 and Cugnon 2012 for reviews of dissenting views and Gründler 2013, Nikolić 2016, Nikolić 2017 for alternative explanations of the Casimir effect.

Although no explicit alternative derivation of the Lamb shift effect has been presented, one can provide a physical interpretation of the result as a side effect of the way matter interacts, instead of electron interactions with the vacuum. As Koberinski forthcoming notes:

What I have been referring to as vacuum polarization effects are at best evidence for fluctuations in the vacuum energy in response to ordinary matter and its interactions. At worst,

¹¹See Koberinski forthcoming for the development of a pretty systematic argumentation against the significance of the CCP – even when it comes to regularisation estimates in Koberinski forthcoming.

¹²Which, he emphasised, was the original intention of Casimir’s work.

one might caution against reading Feynman diagrams too literally; then the vacuum fluctuations as represented in loop diagrams might be interpreted as simply components of nonlinear interaction effects between ordinary matter.

Warning thoughts such as the above seem to be following the spirit of work by Bohr & Rosenfeld (Bohr and Rosenfeld 1933, Bohr and Rosenfeld 1996) on the measurement of quantised electromagnetic field:

we want to emphasize here once again that the consistent interpretability of this formalism is in no way endangered by such paradoxical features of its mathematical representation as the infinite zero-point energy. In particular, this latter paradox, which moreover can be removed by a formal change in the representation that does not influence the physical interpretation, has no direct connection with the problem of measurability of field quantities. [...] A physical measurement of the field energy can be carried out only by means of a suitable mechanical device that would make it possible to separate the electromagnetic fields in a given region from the rest of the field, so that the energy contained in the region could be measured subsequently by application of the conservation law.

Rugh and Zinkernagel 2002 (p. 683-4) insist that, consistent with above, even standard accounts for of effects such as the Lamb shift, Casimir force and the electron's anomalous magnetic moment, which implicate the vacuum, cannot eliminate the possibility that they are simply the side-effects of the way material objects (the plates, the measuring apparatus and the electron) couple with one another.

When it comes to bubble diagrams themselves, however, we face a dilemma. Normally, bubble diagrams are accorded no physical meaning in standard treatments of QFT since they do not contribute to scattering amplitudes of processes such as the scattering of an electron off a proton. So, why include them in the calculation when we are perfectly aware that this is bound to lead to the CCP "disaster"? Perhaps the "sane" option would be to view this predictive failure as a red flag against according bubble diagrams any physical meaning whatsoever. Unfortunately, taking this way out goes against our most deeply entrenched credo about the universal character of gravity; no form of energy can escape the "pull" of the gravitational interaction.

One way to escape the the dilemma in the latter direction is to go further down the skeptical road and exempt, in a somewhat ad hoc manner, vacuum energy from the scope of the gravitational interaction. This essentially amounts to a rejection of the (strong) equivalence principle

and a fortiori assumption (P6). This blocks the transition from the form of the cosmological terms in SR to its form in GR via the standard rule “turn your comas into semi colons and the η -s to g -s”:

$$T_{\mu\nu} \approx -\rho\eta_{\mu\nu} \longrightarrow T_{\mu\nu} \approx -\rho g_{\mu\nu} \quad (4.17)$$

In fact, what the sceptic needs is not so much an argument against the form of $T_{\mu\nu}$ in a non-Minkowski setting as an argument against its very definability or, equivalently, against the very cogency of the semiclassical treatment. Indeed, there are a few points to contest here:

1. **Semiclassical Gravity:** what is the justification for coupling the (classical) GR metric to the average value of the (quantum) stress-energy tensor; while this might be the easiest way to form a classical quantity to enter into EFEs, there is no justification as to why this is the right way forward
2. **Energy Condition Violations:** energy conditions that rule out some pathogenies or exotic physics (such as closed timelike curves or naked singularities) are often violated even before quantum weirdness enters the picture (see Curiel 2017)
3. **Classical Stress-energy Tensor:** there are ambiguities in the way one derives the stress energy tensor for spin-2 theories (see footnote 4, this chapter) that do not necessarily agree with one another on the final result
4. **Definability of QFT:** as we will see soon QFT tools might not be definable outside the Minkowski spacetime structure (because, for example, wildly different spacetimes might require wildly different regularisation schemes)

All this indicates that a (major) revision to the QFT framework might be imminent – perhaps one in which the vacuum does not play a significant role¹³ A similar way out of problem was envisioned early on by Schwinger 1988 whose suggested **source theory** sought to do away with the ill-defined and problematic notion of a non-zero energy vacuum and quantum fields acting on it. In their stead, Schwinger’s theory utilises classical or c-fields and transition amplitudes “sourced” by currents similar to those used in classical field theory. The creation of particles

¹³In fact, Jaffe himself remarks that while his argument undermines the inference from the Casimir effect to the reality of the vacuum, it does not challenge the reality of the vacuum per se. The reason for this being that “Casimir’s original goal was to compute the van der Waal’s force between polarizable molecules at separations so large that relativistic (retardation) effects are essential” but then “following a suggestion by Bohr [25], showed that the Casimir-Polder results could be derived more simply by comparing the zero-point energy of the electromagnetic field in the presence of the molecules with its vacuum values [26]”.

in this theory, what Schwinger dubs “creation acts” is treated in terms of collisions. This gives the theory a more phenomenological character. Critics point out this way out of the CCP is not without its problems. Insisting on technically demanding job of re-expressing QFT in a new language, Saunders 1993, for instance, claims, that it not clear that Schwinger’s approach does away with the vacuum (p. 335):

The remaining method used by Schwinger, whereby the radial component of the stress-energy tensor for the electromagnetic field is calculated, likewise involves an infinite subtraction. It is justified with the words: “No physical meaning can be ascribed to such a term, however, since no stress can arise from a homogeneous dielectric (as such it can be canceled by a contact term)”. This term too has the same form as the expression for the vacuum energy arising in Casimir’s calculation.

to conclude that

Whatever the virtues of Schwinger’s source theory, transparency, and statements of clear and systematic principles, are not among them. I do not believe his methods deliver an unambiguous verdict on this matter.

More work would be needed before the received formulation of QFT were to be jettisoned in favour of Schwinger’s more mathematically-minded alternative.

A more recent alternative meant to phase out the dreaded vacuum state is to be found in the work of Hollands and Wald 2008, Hollands and Wald 2015 which formulates QFT in “a local and covariant manner in terms of locally measurable field observables”. By according operator product expansions of quantum fields (OPEs) center stage, Hollands & Wald “prepare” QFT for a transition to curved spatiotemporal backgrounds eliminating the need for a Poincaré invariant (vacuum) state along the way. They argue that that the fine-tuning involved in the specification of the cosmological constant term is associated with non-perturbative effects, which, compared to perturbative ones, can remain small even when the scales of the theory are high.

While this project is not yet complete and questions the received wisdom about QFT, it also boasts some clear advantages. By eliminating those parts of the theory that anchor it to a Minkowski background, it facilitates the transition to more general spacetimes. An equally strong motivation in favour of approaches that seek to undermine the special status of the vacuum in QFT comes from results such as the Unruh effect. The Unruh effect, derived in treatments of QFT on curved spacetime (Birrell and Davies 1984, chapter 3, Carroll 2019, section

9.4) signals a certain loss of well-definiteness for the vacuum state. In a nutshell, this is because observers accelerating with respect to a given vacuum state will observe thermal radiation contrary to an observer moving inertially with respect to it. The heat bath will have a temperature given by:

$$T = \frac{\hbar}{2\pi ck_B} a \quad (4.18)$$

where a is the acceleration of the observer. Of course, even if this speaks against a global definition of the vacuum, it still is true that locally spacetime is Minkowskian and therefore the rules QFT, the way as we have to know and love them so far, still hold. Presumably, this local contributions of will add up to a global, large scale energy contribution. I am not sure that this is enough to guarantee that something like the CCP arises at large distances, however. As Bianchi and Rovelli 2010 note “the physical effect of the cosmological constant is not visible in this approximation either: it requires to go to very large distance, which is precisely where such local approximation fails”. So, although our tools only suffice to tell us something about the CC locally, *we choose* to, in all likelihood illegitimately, extrapolate to cosmic scales and then ask where things went wrong.

4.2.4 The (Real) Problem

So, what is the real problem with the cosmological constant? Ultimately, I think, we can agree that much like in any “anomaly” the burden does not fall exclusively on the shoulders of any of the protagonists, that is GR or QFT, but rather in attempts to combine them. The whole controversy can be traced back to the complications of putting QFT on a curved background. The QFT framework we currently possess operates on a restrictive set of assumptions that are hard to square with spatiotemporal structures of a more general kind. Examples of such assumptions include the symmetries of Minkowski spacetime, a well-defined and unique vacuum state, unambiguous definitions of conserved quantities etc. Failure of any or all of these conditions indicates not only the need for a substantial revision of the present QFT framework, so as to render it adaptable to a non-Minkowski setting, but equally exposes some key presuppositions for its applicability. These presuppositions also act a precondition for the validity of the EFT framework. This is what we must turn to now.

4.3 EFTs under Pressure...again!

The above discussion is old-school in that it does not utilise the modern EFT methods. Reformulating the CCP using the EFT framework forcefully underscores its commonalities with the violation of naturalness in the hierarchy problem. Once more, the most pressing formulation of the CCP directly links the Λ term to its behaviour under RG transformations:

CCP: The EFT Formulation The cosmological constant parameter Λ corresponds to a relevant operator; its contribution will thus become increasingly more dominant as we go to lower energies. This contradicts our experience, however: the cosmological constant is estimated at a very low number of the order 10^{-50}GeV .

Much like in the case of the Higgs boson mass, the cosmological constant term exhibits a highly atypical behaviour: to obtain results consistent with cosmological perturbations one has to fix the constant at high energies to a very precise value so that it cancels out with the quantum corrections to deliver the value we measure. Apart from the “annoyance” of fine-tuning itself, one cannot help but raise concerns about decoupling similar to those we had to confront in the Higgs case. Once again, the EFT framework appears to be under attack! With a crucial difference: since, contrary to the hierarchy problem, the CCP is not a problem internal to QFT, but involves gravity as well, it points us towards something physically more tangible than a possible failure of reductionism. And since the problem lies in the forced marriage between QFT and GR, the CCP is an indication of the limitations of EFT techniques in non-Minkowskian spacetimes.

4.3.1 Reactions

Since the CCP is formulated as (yet) another naturalness violation problem, the space of solutions or dis-solutions of the problem can thus be similarly characterised:

1. **New Physics:** QFT or GR is to be supplanted by a new theory, which will preserve naturalness. Examples of such extensions include supersymmetry (again!), modified theories of gravity, extra-dimensions and string theory..
2. **Denial:** Denying the problem or downplaying its significance is always a possibility. What involves require in the case of the CCP, however, is a) a reconsideration of the

foundations of QFT (so that one does not rely on concepts such as the vacuum, its energy etc) and b) a re-interpretation of the Λ term in within GR. (As we saw, fine-tuning arguments by themselves will fail to be persuasive if no clear physical meaning is attached to them.).

3. **Failure of EFT framework:** Again, one can avoid the drastic move of abandoning or replacing the EFT framework tout court by restricting the domain of its applicability. The CCP is a valuable guide in this direction.

We will simply omit discussions of the first two alternatives here. The reason is that they closely mimic the corresponding options for the hierarchy problem, which were thoroughly analysed in chapter 3. To avoid reiterating the same arguments here, we will only revisit the third option to update the discussion over the domain of applicability of EFTs with the valuable insights we have gathered through our analysis of the CCP.

4.3.2 Preconditions for EFTs

Our examination of the hierarchy problem revealed that decoupling in some forms, such as that encoded by the decoupling theorem by Appelquist & Carrazzone, does not always hold. Recall that, according to this theorem, whenever we have a renormalisable theory describing the interaction of systems with a scale separation, it is possible to include the high-energy effects into the low-energy description with an appropriate modification to the latter's parameters. We saw that one (physical) way to think of the theorem's failure in the case of the Higgs boson mass is as a signal for UV/IR mixing. This in turn gave us pause as to the range of applicability of EFT techniques. Similarly, in the context of the CCP, the fine-tuning required to make the value of the cosmological constant as evaluated by QFT square with the value required in cosmology also forces a reconsideration of the effective programme on us. The preceding analysis of the assumptions and possible ways out of the the problem have made more salient some of the presuppositions that need to be in effect for an EFT description of a certain physical system to be possible. In fact, the ills of forcing a marriage between QFT and GR shows that the Minkowski spatiotemporal structure might possess indispensable features.

Conditions

1. **Separation of scales:** clearly, there is no sense in which we can talk of an effective theory if we are deprived of a well-defined notion of scales; for you to be higher, I need to be lower. Clearly, this is the point that comes under stress with the hierarchy and cosmological constant problems – but also with the possibility of a UV/IR correspondence.
2. **“Stability” of background structure**¹⁴: one needs a well-behaved spatiotemporal structure, i.e. one that is typically treated as a weakly perturbed flat space. This can be encoded in the demand for **uniqueness of vacuum**: to apply the standard QFT tools one typically needs an unambiguously defined (that is, invariant under any spatiotemporal transformations) vacuum state on which the field operators act.

- (a) **Bounded Hamiltonian:** a more relaxed condition is that energy is bounded from below to avoid a “turtles all the way down” problem similar to that Dirac’s equation was facing in its early days of conception (and for which the Dirac sea idea had to be introduced)
- (b) **Killing vector field:** the uniqueness of the vacuum can also be derived from a globally defined killing vector field for the metric $g^{\mu\nu}$. This will allow the definition of some conserved matter energy over spacelike slices:

$$H_{mat} = \int_t d\Sigma_\mu K^\nu T^\mu{}_\nu \quad (4.19)$$

with $d\Sigma$ some measure over the slice.

- (c) **Slow-evolution:** the background must only be evolving slightly with time, i.e. the process should be adiabatic enough, to ensure that no heavy degrees of freedom will be produced. It is also important to prohibit the evolution from low-energy to high-energy states (according to a set cut-off) and vice versa. This will make degrees of freedom suddenly become relevant or irrelevant.
3. **Locality:** the uncertainty relations guarantee that high energy processes exceeding the energy cut-off will only last for infinitesimal time intervals and no locality violating propagation will take place. Entanglement (say, between high and low degrees of freedom) might complicate things.

¹⁴A more detailed and technical discussion of these points can be found in section 3.3. of Burgess 2004

4. **Naturalness:** one might add here the demand that parameters at a given scale M do not receive contributions from higher scales or any other form of naturalness constraint deemed important such as the absence of fine-tuning

Assessment The natural question to ask now is: what evidence or hints of new physics do we possess that any, or even all, of these conditions fail to hold? Although mostly based on tentative results or some speculative suggestions, there are indeed indications from cosmology, the foundations of QFT in curved spacetime and even string theory that some of these conditions will probably need to be forsaken at some point in the future:

UV/IR Correspondence : there’s a growing interest in the way gravity seems to efface the lines between high and low energy degrees of freedom. An important case in point is black hole physics for the fact that “black holes radiate at temperatures inversely proportional to their masses necessitates some sort of ‘UV/IR mixing’ in gravity — infrared physics must somehow ‘know about’ heavy mass scales in violation of a naïve application of decoupling” (Craig and Koren 2020, p. 3). It is probably not that surprising to learn that gravity and physics at the Planck scale behaves in ways that are alien to our currently prevalent conception of doing physics. Yet, the mixing of UV and IR scales is a violation of a key EFT assumption about the separation of scales. It is unclear what bearing this should have to the hierarchy problem (a point to which we shall briefly return in the next section), but it does imply the need for a deeper (re)thinking of the decoupling assumption – and naturalness along with it.

Possible indication in favour of UV/IR mixing also comes from some low-energy limit models of open string theory associated with non-commutative field theory (NCFT). NCFTs begin with the introduction of a nonzero commutator between position operators:

$$[\hat{x}_i, \hat{x}_j] = i\theta_{\mu\nu} \tag{4.20}$$

which leads to the breaking of Lorentz invariance. The main complication with treating such models within the EFT framework is that one has to define a new product for fields, the so-called Groenewold-Moyal product, which fails to be local because the infinite number of derivative operators it contains. Unsurprisingly, this also goes against the locality assumption as well because a key EFT idea is that for interactions happening at a local level (around a small neighbourhood of a point), one is allowed to “collapse” or “squeeze” the propagating

underlying degrees of freedom to point-like interactions.

—→*rethink: conditions 1, 3, 4*

Cosmology & Swampland : Ideas such as emergent universes provide another possible instance of failure for the EFTs. As discussed by Brandenberger 2021, the emergent universes is an alternative to inflation in which the expansion of the Big Bang cosmological model does not arise from an inflaton field, but rather an early phase that is cannot be treated by standard EFT methods. Brandenberger (2021b) also emphasises that standard EFT techniques might be incompatible with models of early universe cosmology models. He traces this incompatibility to the exponential increase in the wavelength of cosmological fluctuations for:

if the phase of inflation lasts sufficiently long, then length scales corresponding to those which are currently measured in CMB anisotropy experiments and large-scale structure surveys originate at the beginning of the period of inflation with a physical length smaller than the Planck length.

To remedy the problem, the “Trans-Planckian censorship conjecture” has been proposed as an additional constraint on EFTs of the early universe. This principle essentially imposes a constraint on the wavelength of fluctuation modes so that they may not grow at a super-Hubble scale. In general, a common theme in the examination of the connection between low-energy physics and the more fundamental high-energy theory, string theory, to impose constraints or conditions in the form of conjectures or principles for the kind of features that an effective theory should have in order to be completed into quantum gravity (and thus be part of the so-called “string landscape”) or not (and thus be part of the “swampland”) (Graña and Herráez 2021).

Other problems of the same vein have to do with possible violations of Lorentz invariance either through the introduction of some minimal length or via the modification of Lorentz invariance itself, as is the case in DSR (Doubly Special Relativity) models. Doubly Special Relativity is “doubly” special because it imposes not only an invariant velocity scale c corresponding to the speed of light, but also a fundamental length (or inverse-momentum) scale ℓ_{DSR} on which inertial observers must agree. This leads to a modification of the dispersion relation like:

$$E^2 - p^2 c^2 + \ell E p^2 c - m^2 c^4 = 0 \quad (4.21)$$

In a radical departure from common wisdom, this also leads to transformations that are trans-

lation dependent so that interactions that might appear local for one observer might appear non-local for another translated according to the DSR transformations (Amelino-Camelia et al. 2011)! Evidently, this renders the applicability of EFT techniques, which are grounded in the local nature of interactions, a quite non-trivial matter.

—→*rethink: conditions 1, 3*

Entanglement : Cosmology opens up yet another avenue for the investigation of the possibility of “mixing” of degrees of freedom across separate scales. This is because the evolution of the early universe, as described by inflation theory at least –with quantum fluctuations being “stretched” along with the rapid expansion of the universe– provides a unique context in which small and large scales influence one another. In this context, it has been suggested that entanglement, the quantum feature par excellence, is responsible for a blending of the UV and IR that invalidates the standard conception of Wilsonian EFTs (Brahma, Alaryani, and Brandenberger 2020):

The crucial thing to note in this case is that the standard Wilsonian effective action does not exist since the sub-Hubble modes, which are integrated out, are not excluded by any conservation law [8, 9]. This is contrary to traditional EFTs in which energy conservation ensures that high-energy dofs cannot be part of the system if they were not initially present, and the entire dynamics of the low-energy system can be described by an effective Lagrangian consisting only of the light dofs.

This leads Brandenberger 2021 to turn to a consideration of open EFTs in which (p. 3):

...one studies the (non-unitary) evolution of the reduced density matrix obtained by tracing over the unobservable dofs. These methods are prevalent in other branches of physics, and we adapt them here for cosmology, in the hope of measuring the degree of entanglement, between our super-Hubble system modes with its environment (see [16,17] for entanglement entropy calculations in this setup), through observable predictions.

Open EFTs suggest a modification to the standard treatment of the subject where the usual split between UV and IR degrees of freedom can be a liability rather an asset. For example, Agon et al. 2018 mentions the AdS/CFT correspondence as a specific case where one would find the split undesirable - since it is the connection between UV and IR that one would want to preserve.

—→*rethink: conditions 1, 3*

QFT in Curved Spacetime : No need to excuse into exotic lands, however. The more down-to-earth project of transplanting QFT from its Minkowski background to more general spacetimes also comes with significant difficulties that violate some of the preceding conditions. First of all, as has been frequently noted (e.g. Burgess 2007), the very existence of an evolving, “non-stable” (i.e. non-adiabatically evolving) background structure creates complications with level-crossing (production of heavy particles) or when there is no neat separation between negative and positive energy frequencies (Birrell and Davies 1984, ch. 3.2).

Furthermore, the very mathematical-conceptual framework used in standard treatments of QFT might need substantial revisions. As noted by Hollands and Wald 2015 in their approach to this task, while assumptions such as: i) the distributional nature of quantum fields, ii) their local and covariant construction and iii) the satisfaction of conditions on the energy spectrum can be transferred to the curved spacetime setting, the existence of a preferred Poincaré-invariant state, the vacuum, cannot be guaranteed in this new setting. In fact, Wald 2018 goes as far as to compare the existence of a unique CS-QFT vacuum to a preferred coordinate system in GR:

After our more than 90 years of experience with classical general relativity, there is a consensus that it is fruitless to seek a preferred coordinate system for general spacetimes, and that the theory is best formulated geometrically, wherein one does not have to specify a choice of coordinate system in order to formulate the theory. Similarly, after our more than 40 years of experience with quantum field theory in curved spacetime, it seems similarly clear to me that it is fruitless to seek a preferred vacuum state for general spacetimes, and that the theory should be formulated in a manner that does not require one to specify a choice specify a choice of state (or representation) to define the theory.

It is thus to be expected that the vacuum, forever a peculiarly elusive concept in QFT, could face threat of extinction in the curved spacetime context. Yet, the repercussions of such a move are pretty dramatic. For it undermines the very foundation of the conceptual framework of doing QFT. Both standard approaches, canonical quantisation and the functional integral approach presuppose the existence of a ground state (either free or interacting) and proceed to construct fields as mode expansions over these states. Alternative approaches, such as those based on algebraic structures, can do away with the vacuum state and are thus better suited for an extension to a non-Minkowski setting. In fact, the approach developed by Hollands and Wald 2015, taking the operator product expansion [OPE] as the central concept, starts from an algebraic formulation of QFT to extract the key ingredients for extending the framework

to curved spacetimes. The OPE transforms a product of local operators $\mathcal{O}_1, \mathcal{O}_2$ evaluated at different points x, y as an expansion over composite local operators as the points come close $x \rightarrow y$:

$$\lim_{x \rightarrow y} \mathcal{O}_1(x) \mathcal{O}_2(y) = \sum_n C_n(x-y) \mathcal{O}_n(x) \quad (4.22)$$

with C_n called the Wilson coefficients are independent of the external states that mathematically are to be described as distributions. Hollands & Wald (2010) claim that the existence of such a product whose coefficients respect a set of conditions¹⁵ can act as a potential replacement of the Poincaré invariant state (vacuum). In this way one does with the very ambiguous element of QFT in curved spacetime.

It is unclear what the role of EFTs will be in such a new context as more work needs to be done in the direction of interacting fields and renormalisation. However, it is reasonable to expect that certain weaker conditions (weaker, that is, than the existence of a unique vacuum state) will suffice to guarantee the applicability of the EFT techniques. Burgess (2007), for example, claims that conditions such as the adiabatic evolution of background structure or the mere existence of low-energy observers suffice to ensure that EFTs will apply even in worrisome contexts such as black holes and inflation. Once again, however, this is possible only as long as specific conditions on the targeted models are imposed (such as that high-energy modes enter low-energy physics in their ground state) - consistent with the pre-conditions discussed above.

—→*rethink: condition 2*

4.4 Conclusions

The hierarchy problem gave us the first indication that something might be wrong with the EFT framework. Depending on how much one wants to read into naturalness, its violation indicates the failure of some form of decoupling. If one equates or pegs this to reductionism, then naturalness can assume the role of a precondition for the applicability of EFTs. I take it that the previous and the present chapters have showed that this is not a necessary step. Starting from the different reactions to Wallace’s argument we were led to the issue of uncovering the presuppositions of the EFT framework. Our excursion into GR via the cosmological constant

¹⁵For the curious, conditions imposed on the C_n -s include locality, covariance, the existence of an identity element, commutativity or anticommutativity (depending on the respective fields), associativity, analytic dependence on the metric and more technical such as asymptotic positivity, compatibility with the introduced operators etc.

problem and the disputed territories of QFT on curved spacetime have allowed us to derive a compendium of conditions that must hold to avoid a breakdown of the EFT techniques. Seen through this lens naturalness can be seen as yet another stricter-than-decoupling condition that one can impose on the parameters of an EFT. Therefore, the whole discussion in this chapter has highlighted the flexibility of the EFT framework: when investigating BSM physics one can form a huge space of possible theories that can be written down consistent with some (chosen) fundamental principles (which are nevertheless themselves revisable) and then appropriately constrain them with specific conditions (additional principles) or desiderata of a successor theory. This is reminiscent of the way we treated the construction and modification of models and theories in chapter 1. A given model-theory is “tailored” to guarantee that some necessary features will be preserved into the unknown region ahead.

Chapter 4b

Interlude: The World beyond EFTs

Once more, it is time to take stock. In our previous interlude we considered several possible stances one can take with respect to the EFT framework, ranging from full-blown instrumentalism to excessively optimistic realism. In the meantime, however, we have examined two central problems in contemporary theoretical physics and their relation to effective theories. We saw that when the problem in each case is transcribed into the language of EFTs, it casts into doubt the straightforward application of the RG techniques and the EFT reading typically attributed to them. It is about time we turned to our (final) assessment of the four stances one can take with respect to EFTs. Recall that these are:

1. **Scepticism:** question the validity of EFT reading of QFT; opt for a more rigorous mathematical basis that is found in axiomatic projects like AQFT, constructive QFT etc
2. **Caution:** the EFT framework has rightfully shaped our understanding of QFT in the past decades, but there might be reason to restrict its scope
3. **Optimism:** EFTs constitute the right framework in which to think about QFT; since they instantiate a certain conception of how to do physics, which pervades the whole field, they should be taken to apply universally
4. **Enthusiasm:** adopt the most radical lessons from the EFT perspective such as an abandonment of the project of a “final theory”, reductionism and foundationalism

4b.1 Appealing the Verdict

In our latest “interlude” we left the battlefield just as the optimist camp was steadily gaining the upper hand. Recall that after the first few skirmishes they had made advancements across the whole front: their understanding of QFT as an EFT was found to be both plausible and promising, their brand of realism, effective realism, suffered some losses but its core tenets were for the most part left unscathed and, mostly importantly, the timidity of their more conservative opponents left them with little room to manoeuvre at this first stage of the battle. In light of this, our preliminary verdict was in favour of the optimist’s reading of the EFT framework.

Our long campaigns in unnatural lands, however, have allowed us glimpses into novel mysteries that have shaken our confidence in the optimist narrative. In particular, violations of naturalness in the SM as well as cosmology have illustrated the potential shortcomings of the EFT framework in a striking manner. From a complete failure of reductionism to a restriction of the EFT framework or the seemingly innocuous reappraisal of the physical significance of naturalness, no matter what choice ones makes, the final product is a further constraint on the EFT framework. Here is why:

- ⇒ Failure of **reductionism**: hardly needing any explaining, if reductionism fails, the prospects for an effective realism are shrinking – this is too a radical a breakdown of our autonomy of scales expectations.
- ⇒ Failure of **decoupling**: if naturalness, construed as a form of decoupling, is abandoned, QFT must be seen as a singular case and effective realism can best be said to apply within a narrower regime
- ⇒ Failure of **naturalness**: if naturalness is not equated with decoupling, then a whole lot more of the optimist view can be salvaged, albeit with some cost: we switch to a more heuristic or informal notion of “effectiveness”.
- ⇒ Failure of **the standard model**: this is the most popular view in the physics community and the one that requires no concession on the part of the effective realist. Unfortunately, it has come under a severe amount of pressure from empirical evidence.

Given the impasse reached with recent runs of the LHC, the only options left for us are not wholly favourable to the optimist position. A fortiori, the enthusiast’s position, an over-ambitious extension of the optimist’s, is all but ruled out. The question we must turn to instead

is this: which of the most conservative alternative stances is best warranted in light of our analysis?

4b.2 A (Final) Verdict

Recall that the pessimist wants us to abandon the EFT reading of QFT and instead opt for a mathematically rigorous and consistent QFT for which asymptotic expansions will only be approximation tools and not an essential feature of theoretical representation. If one does not want to go all the way down this path, the cautious stance is the sole position consistent with the analysis we have made so far. So far our presentation has defined the position in a negative manner, i.e. through a rejection and elimination of the alternative positions. However, one can give a more positive content to this position through the following basic tenets:

The Cautious Temperament

- Effective realism encompasses a “philosophy” about the way we do physics: separating scales, splitting up the degrees of freedom and focusing only on those relevant for the scale under investigation. The ontological picture is pretty similar to that of the optimist: nature largely allows for a separation of scales so that we can provide accurate descriptions (up to a margin of error) at one level without knowing the physics at a more fundamental one.
- The EFT framework adds a stricter requirement on physical theories by making decoupling a much more formal procedure via the RG, the behaviour of parameters across scales, naturalness-like constraints and so on. These tools allow us to better understand the “approximate truth” of our theories. The manifestation of processes (and the appearance of structures) becomes relevant only when the right scales have been triggered.
- Nevertheless: while the expectation of decoupling in the less stricter sense of “details about the more fundamental state of a system do not matter” is corroborated, one needs to be vigilant about possible breakdowns of the framework. These may originate in the failure of certain preconditions to hold. Such conditions may include:
 1. The existence of a unique vacuum state or at least a way to distinguish between positive and negative energies (complicated in curved spacetimes)

2. Also relevant to above is the existence of stable (slowly evolving) background structure
3. A neat separation of scales, so that it is meaningful to talk of higher order and lower order processes. (but: UV/IR correspondence)
4. Locality – absence of entanglement between macroscopic and microscopic degrees of freedom (as in inflation). Otherwise, tools like open EFTs or a whole new theory might be required.

□ Failures of the EFT framework signal its restrictiveness as a framework and should not be equated with a more dramatic breakdown of the way physics has been conducted. Delineation of appropriate regimes of applicability exorcises the phantom of transcendentalism. This, in effect, implies that the tools and strategies that have been employed in particle physics throughout the 20th century (assuming that naturalness is one of them) may simply not apply to regimes such as early universe cosmology or might require further refinement in the case of broken symmetries – especially when it comes to the dynamical origin of phase transitions.

□ In the final analysis: all theories are effective in a loose sense, but far fewer fulfil the requirements in the much stricter sense needed for the EFT frame. Violations of naturalness as in the cosmological constant or Higgs case reflect this latter requirement and can be seen at worst as restrictions of the framework.

Yes, but... We must retain our optimistic outlook on EFTs and post 70s QFT, but at the same time appreciate the conceptual limitations of this framework, especially given the lack of empirical support in favour of BSM suggestions which were largely motivated by naturalness. Appreciating the lessons we have been taught by this change in perspective in fundamental physics is most certainly the way forward for philosophers studying the foundations of QFT, but as is frequently the case in life: caution is advised.

Chapter 5

Coda: Philosophical Lessons

The above investigations lead to some interesting prospects for our understanding of EFTs within the broader edifice of physical science and its philosophy. Nonetheless, there still remains room for a more exhaustive treatment of some issues of particular importance in contemporary philosophy of science. Our insight into the inner workings and foundations of EFTs can shed light to both ontological and methodological investigations. For the former, the “refutation of fundamentalism” entails a revision of the role of reduction in thinking about the structure of the world. As for the latter, the EFT-motivated methodological version of realism as well as its potential limitations (along with the controversial role played by naturalness) lead to some interesting lessons about theory construction and the role of extra-empirical confirmation principles.

5.1 Ontology: Realism and Reduction

The two Rs –Realism and Reductionism– have come to dominate contemporary philosophical thinking for the past 40 years. Reductionism has been the treasured child of particle physics throughout the 20th century. Disavowing first what critics of high energy physics’ (almost monopolistic) voracious appetite for governmental funding have dubbed:

- **theory reductionism**, i.e. the idea (or dogma?) that all special sciences will eventually be “absorbed” into physics so that descriptions of phenomena across all scales will come in terms of the fundamental posits of our (final) theory of physics
- **methodological reductionism**, i.e. the idea that in order to make progress in higher-level science some progress in lower-level sciences is required (think of psychology and

neuroscience)

Weinberg 1987, in his plea for the construction of a larger particle collider, sought to justify the primacy of fields such as high energy physics by appealing to a more modest “doctrine” of:

- **objective reductionism**, i.e. the idea that the “explanatory arrows” in science go from lower (in terms of depth) to higher levels and emanate (or at least seem to) from a common origin

Thus, physics, and in particular particle physics, enjoys primacy over other scientific fields (or subfields of physics itself) in virtue of being the most “fundamental” discipline. Fundamental in that despite the fact that laws at any level of physics may reveal some novel or “emergent” character, (almost) no one seriously doubts that this is ultimately due to the nature and the configurations of their microscopic constituents as revealed to us by (fundamental) physics.

It is not so clear how this rather loose basis, as Weinberg himself seems to admit, can help support the additional siphoning of resources towards particle physics. In fact, the heart of the problem does not even concern the mere practicality or possibility of deriving the (more) macroscopic physical, chemical or biological, psychological and social phenomena from fundamental physics, but rather a deeper conceptual issue with the explanatory value of a project whose motivation stems from a desire to find *the* ultimate unified theoretical framework. This aspiration is nicely captured by Einstein (1918 – quote in Howard and Giovanelli 2019)

The supreme task of the physicist is ... the search for those most general, elementary laws from which the world picture is to be obtained through pure deduction. No logical path leads to these elementary laws; it is instead just the intuition that rests on an empathic understanding of experience. In this state of methodological uncertainty one can think that arbitrarily many, in themselves equally justified systems of theoretical principles were possible; and this opinion is, in principle, certainly correct. But the development of physics has shown that of all the conceivable theoretical constructions a single one has, at any given time, proved itself unconditionally superior to all others.

Our explorations in the structure of theories and EFTs led us to a view that deems the project discussed above, if not obsolete, too restrictive. What we realised throughout our investigations is that theories latch onto the world in such complex ways that the whole project of a derivation of physical phenomena from a set of primitive laws (or “given-s”) is ultimately rendered moot. The purist’s reductionist agenda of constructing a neatly compartmentalised pyramid of

knowledge, where each level is derived from or reduced to the previous one, sounds more like a theoretician's dream rather than a methodological stance that can survive contact with actual practice. We saw that EFTs employ forms of reasoning and model construction analogous to those employed in fields like condensed matter physics, statistical and continuum mechanics in which "purist-leaning" bottom-up derivations are combined with top-down constraints as well as empirical input. Models of this sort do enhance our understanding of the world and help us form expectations about experiments and future phenomena. Whether they can be also seen as explanations in a stricter philosophical sense really hinges upon one's preferred theory of explanation¹

Therefore, the more situated, pragmatically-oriented and strategising aspects of theories coupled with the more formal aspects one is more familiar with from textbooks are consistent with a moderate version of reduction. Perhaps this can also be viewed as a reduction of a methodological kin, which (rightly) emphasises the potential insights to be obtained as we probe further into the constituents of systems but (crucially) parts ways with the standard reductionist thesis once this seeks to establish itself as an essential requirement for scientific theorising. For, on the one hand, elevating the hope that the huge mosaic of intertwined of horizontal and vertical theoretical frameworks converges to some set of fundamental principles to the status of some transcendental principle frequently founders off scientific practice. On the other hand, construed as a purely empirically justified project, reductionism in this stricter sense appears to be undermined by the EFT-inspired methodological points we have made.

Most crucially, it must be stressed that this revisionary account of the reductionist agenda does not entail an abandonment of our realist aspirations. We take it that scientific theories, even when not fully rigourised, universally encompassing or completely compatible with one another may still aspire to represent the world – albeit in highly non-trivial and vastly intricate ways. On the contrary, loosening the reductionist claims might as well be seen as shielding true scientific theorising against unnecessarily strong demands. The more relaxed demands of the

¹It would not hurt to digress on this topic a bit. The singular nature of RG explanation has attracted the attention of much work in recent philosophy of science. Starting roughly with the work of Batterman 2001 on asymptotic reasoning, philosophers have sought after alternative accounts of explanation beyond the established D-N model (Hempel 1965), causal (e.g. Salmon 1984) or unificatory accounts (Friedman 1974, Kitcher 1989). The problem with causal accounts, as argued by Batterman (Ibid) for example, is that the universal behaviour exhibited by asymptotic phenomena of radically different fundamental constituents belies the central role accorded to microphysical descriptions. In other words, the explanation of the fact that systems so disparate as interacting quantum fields, spin lattices, fluids, solids and so on are accurately described using the same RG equations cannot lie in the microscopic configurations of these systems. After all, it is exactly their microscopic constitution that varies so wildly!

view favoured here consist of:

1. **Methodological realism:** motivate a shift to a form of realism that is attentive to the existence of structures within a theory that ought to be read purely instrumentally as opposed to representationally. It takes as the distinctive mark of representing structures their “confounding” and open-ended character. Ontological claims are made in the light of historical and methodological analyses.
2. **Contextualised vocabulary:** in proximity with Putnam’s point about the indispensability of higher-order explanations, the new framework liberates us from the shackles of bottom-up reasoning, allowing for theoretical vocabulary to be adjusted to a specific range of phenomena (quarks - atoms - particles, q.fields - wavefunctions - c.field etc) as suited.
3. **Epistemic modesty:** the methodological intricacies involved in analysing a received theoretical framework and constructing a new one imply that we cannot read-off the ontology of a theory in a straightforward manner, but must rather pay attention to the way a theory is applied, the experimental techniques used to test it, no-go theorems etc.

Indeed, resistance to the above claim might come from some residual commitment to what we have earlier on alluded to as “fundamentalism”. We suggested three different guises, relevant to our purposes, that this attitude might take:

1. **The Meta-physicists:** claim that a theory needs to be understood from the ground up as making claims about the structure of the world at its most fundamental level
slogan: “not ready to take lessons in ontology from quantum physics as it now is!”
2. **The Classicists:** insist that a well-defined theory must offer a clear-cut ontology and a dynamics governing its evolution in space and time
slogan: “define your (primitive) ontology!”
3. **The Literalists:** relax the above requirement but favour “straightforward” interpretations of a theory in light of its canonical (i.e. as standardly practised) formulation
slogan: “take your theory seriously!”

All of these stances come under pressure once we sever the injective mapping between theory and reality. For most clearly, no one taking the approximate, effective character of physical

theories seriously and espousing some form of naturalism in metaphysics would contemplate the idea that we can penetrate into the “fabric of reality” at its most fundamental level in one fell swoop. The story of physics, with sufficient hindsight², is a story of progressive vertical layering and horizontal patching of theories towards an ever(?) fleeting reality. Similarly, the “classicist”, whose inspiration stems from the “canonical” picture of classical mechanics as a theory of particles moving and interacting in space also seems to run on a retrogressive view of scientific theories. When writing down an effective Lagrangian one is often working phenomenologically, i.e. making an educated guess about the fields, the kind of processes and the parameters that will play in the description of a given system. This is often done by neglecting or ignoring the underlying physics³.

The “literalist” approach is somewhat trickier to address along the same lines, but I think that it is ultimately reliant on an understanding of theories that is also put under pressure by an EFT reading of scientific theories. Why? Because of an inherent ambivalence in the maxim of *taking seriously* a theory. A straightforward reading of the ontology off the mathematics of a theory and the way its structures come to acquire a representational role is almost always mediated by a rather complex analysis task involving tools from formal theorems (like spin-statistics, Weinberg’s theorem), no-go theorems to experimental facts. Simply reading what is “out there” from the (most) abstract mathematical formalism of a theory is not enough. A more nuanced approach to theory interpretation should be adopted. EFTs may contribute to such an analysis in ways examined in the following section:

5.1.1 Theory Interpretation in Light of EFTs

Acknowledging the effective and approximate character of QFTs and using the technical arsenal of EFTs, one must factor in some further considerations when examining the ontological import of a theory. Pretty schematically, the main steps to an attempted interpretation of a theory are the following:

1. Define the mathematical space.

²Although it is certainly interesting to appreciate the prescience of figures such as Newton whose stance was to treat scientific theories as continuous approximating probes into the deeper nature of reality.

³Again, it is important to emphasise that this does not undermine the merits of such a more stringent ontological approach in foundations of QM (especially the necessity of specifying “local beables”. After all, a lot of the “complaints” in that context are directed against the current operationalist framework adopted. Nevertheless, these remarks do highlight the somewhat passé –and probably dispensable– character of the requirement for explicit ontological commitments in advance and independently of the dynamical features of a theory usually in the form of some “primitive ontology”.

2. Declare relevant structures to be involved in a particular description.
3. Identify the structures carrying explanatory weight.
 - (a) Check for scheme independence.
 - (b) Check for special conditions.
4. Provide appropriate conditions to characterise structures.

Step 1: The first step roughly corresponds to the adoption of a framework, in the manner defined in chapter 1: it determines the kind of mathematical toolbox to be used. For example, in the case of QFT this will involve operator-valued distributions for fields, commutation relations for their compatibility, connections for the gauge bosons etc. Typically, QFTs operate within this framework combining quantum and relativistic principles. EFTs will naturally also “live” in this framework, albeit occasionally necessitating appropriate modifications to be applicable. For instance, we saw that they might impose additional conditions on the background structure presupposed – with curved spacetimes constituting an important challenge to the well-definiteness of the hierarchy of scales.

EFTs play an even more crucial role is in steps 2 and 3. The construction of an EFT not only involves a procedure by which the physicist decides what degrees of freedom will be relevant for the system investigated, but can also be used to deny representational significance to some of the structures invoked:

Step 2: Simply identifying the (mathematical) tools to be used is not enough. We want to see how they are to be used in describing concrete phenomena. At this stage, we identify those structures relevant to descriptions of specific physical processes: from idealized free fields to interacting ones, whether they are spinorial or vectorial, how they couple with one another and so on. For example, in the case of interactions between electric charges, one will need to identify fields ψ for the description of charged matter and the force carries A^μ for the electromagnetic interactions. In this case, we know that the former require spin $\frac{1}{2}$ representations of the Lorentz group, corresponding to fermionic degrees of freedom while the latter, i.e. the photon field, are of spin 1, corresponding to bosons. The interaction is given by $\bar{\psi}A\psi$. Depending on the problem at hand, that is the kind of physical process that they want to describe, physicists will use some structures and/or their properties and simply discard/disregard others that they want to abstract from. This means that if they might wish to

describe, say, how quarks interact qua carriers of electric charge but ignore their colour charge (and thus strong interactions) they will treat them as fermions coupled to the electromagnetic field.

Evidently, effective approaches can be particularly helpful in this task of identifying relevant structures. By focusing on interactions that are relevant up to the energy scale triggered, one greatly reduces the number of graphs that will be required, the number and, crucially, the kind of fields involved as well as avoid complexities accompanying deeper, more encompassing descriptions. One can thus “get away with” representing the proton as a simple particle interacting with an electron (as opposed to treating it more fundamentally as a complex formation of bound quarks) or “forget” about W bosons and the physics of electroweak interactions in the Fermi theory approximations or even hide higher order contributions within the internal lines of interaction diagrams when the energy scales triggered are low.

EFTs play an even more vital role in descriptions of regimes in which the full theory fails to apply. This is the case with Chiral Perturbation Theory, for example. QCD, the more fundamental theory describing the physics of quarks, cannot be used in the standard perturbative manner at the low energies involving pions and mesons. At energies close to 200 MeV quarks and gluons are no longer the right degrees of freedom to track. Instead, the description shift to conglomerates of quark pairs like mesons. One can then treat formations like pions as simple structures in order to track their interactions. The separation of scales, the key insight of the EFT approach, is what enables these descriptive opportunities.

Step 3: Some physical results should not depend on the kind of description one adopts. In mechanics, for instance, we do not want observers, who disagree on their chosen frame of reference, to also disagree about the causal order of events. Similarly, in the context of QFT, results about the frequencies of interactions that we observe in our experiments, should not be dependent on choices such as the regularisation and/or renormalisation schemes employed. Discrepancies between the schemes (the way we come to resolve the ill-defined multiplication of distributions, the parameter scale we set etc) should be mere descriptive artifacts – not representing new physics. This is something explicitly highlighted by the RG, which tracks the way parameters and fields adjust to maintain the physically significant part, such as the correlation functions, invariant. The EFT techniques have led to some amendments to the renormalization schemes. The main issue was to make manifest the decoupling of heavier degrees of freedom

such as muons even at intermediate stages of the renormalisation procedure⁴. To this effect, hybrid methods of renormalisation have been devised so that one can retain a mass-independent scheme like minimal subtraction but at the same time make apparent the degrees of freedom dropping out at each level. This is done by retaining minimal subtraction as the chosen scheme but switching EFTs across scales and performing the matching between the various levels separately and, essentially, “by hand”. This hybrid scheme is called decoupling subtraction (for more details see Burgess 2020, sec. 7.2.3).

Step 4: While EFTs play no distinctive role in this step, some of their insights and presuppositions can be useful guides to ontology. On the one hand, debates over the correct ontological interpretation of QFT (particles vs fields) have for the most part focused on the coherence of a particle ontology and, accordingly, have spent much energy to spelling out what an appropriate particle notion would look like (e.g. Teller 1995, French and D’Eccio Krause 2006). On the other hand, to address worries about over-commitment, it is necessary to provide further conditions that can deny certain mathematical structures representational status. The discussion in section 9.4 of Duncan 2012 is an example of the kind of investigation onto the desired features of ontologically significant structures.

To characterise the particle-like character of fields, classifying the operators appearing in the Lagrangian will not be enough, but will also require something to be said on whether the respective states constitute eigenstates of the momentum operator, their stability at asymptotically large times etc. In fact, given well-known complications in the unambiguous definition of particles in curved spacetime (see Birrell and Davies 1984, chapter 3), accurately exposing the particle-like disposition of fields will probably require additional facts about the way detectors work: how localised excitations are identified, how deflection angles are spotted etc (Martin-Martinez 2015, Grimmer, Torres, and Martin-Martinez 2021, Polo-Goméz, Garay, and Martin-Martinez 2021). These facts will also be useful when we characterise stable structures across different scales. A compound particle such as the proton can look simple at the appropriate energy scale.

Conclusion We accept that the phenomena at a given scale are (at least in principle) reducible to the laws of microscopic physics. Nonetheless, accepting some autonomy between these

⁴The problem is that a mass-independent scheme does not explicitly depend on the mass scale of the degrees of freedom involved and as a result, the contributions of fields with different masses like muons and electrons are not differentiated in the coupling runs.

various levels, we are not necessarily committed to the idea that a) our understanding of a given scale requires a description in terms of the underlying micro-regime and, following from this, b) we remain agnostic as to the possibility (and necessity) of a “final theory”. The EFT framework allows us to systematically patch together theories at separate energy scales so as to guarantee the coherence of a “relaxed” version of reductionism. This is to be construed as the minimal demand for the compatibility of two level descriptions along with a minimal correspondence assumption (which might be called “compositionality principle”): that such and such structures (be they protons, neutrons, nuclei etc) are “born” out of such and such structures (be they quarks, nucleons, strings etc).

5.2 Methodology I: Theory Construction

Apart from metaphysical lessons, the EFT-inspired account of scientific theorising developed has a lot to offer towards refinement of our methodological presuppositions. Quite expectedly, there is both a destructive and a constructive aspect to this. First and foremost, our critical analysis of the hierarchy and cosmological constant problem leads to a reevaluation of the status of certain heuristic principles such as naturalness. A further re-examination of the status of broader demands such as fundamentalism or reductionism and an autonomy of scales must follow suit. Furthermore, the rejection or restriction of such constraining, albeit reasonable, tenets can open up routes of research that have, up to this point, remained largely unexplored either because most productive energy went towards directions compatible with the prevalent constraining principles or because alternatives were considered mostly quixotic – if not full-blown impossibilities.

Perhaps the central constructive aspect of the effective methodological paradigm is the flexibility it allows for in theory construction. Starting from a well-established theory and assuming it to apply only to a certain range of phenomena, i.e. treating it as an effective theory, one can attempt to extend it by adding further terms to the relevant expansion, modifying certain parameters, incorporating further properties of its target system(s) as relevant and so on. An example along this direction is research in programs such as modified theories of gravity where the standard Einstein-Hilbert action in GR is treated as a low-energy approximation to a more complete theory (e.g. Burgess 2020, chapter 10). If GR is to be treated as an EFT one must

expand the original GR Lagrangian to include higher order terms as corrections:

$$\mathcal{L}_{mod} = \mathcal{L}_{GR} + \alpha R^2 + \beta R_{\mu\nu}R + \gamma R_{\mu\nu\lambda\rho}R^{\lambda\rho} + \delta \square R + \dots \quad (5.1)$$

The flexibility ensured by these effective approaches consists in the fact that physicists are free to enter into a zone of speculation, where trying out new ideas (as dramatic as the change of topological features of the underlying geometry as in approaches that replace Riemannian geometry with Finsler geometry) does not beget any form of formal or conceptual justification, but is only vindicated in light of the proposed model's capacity to solve key problems that have been identified as important within the community. The approach is also conservative in that one prefers to cautiously grope in the dark for hints to move forward as opposed to ambitious projects seeking to resolve all problems in the Hegelian way, i.e. re-systematising the full body of our knowledge according to some (new) set of principles. With the quest for a fundamental theory (almost) dead on its track and the undermining of reductionist programs similar to Weinberg's above, we find that progress in physics, for now at least, is largely independent of plunging further into the oceanic rifts of particle physics.

This point is strengthened when we taken into account the dynamic character of scientific theories. The interplay between top-down empirical constraints and the more meticulous bottom-up construction of models blares the neat separation between the heuristic and the rigorous in favour of a more situated, trial-error and strategising approach to theory construction. Consistent with the techniques we described in chapter 1, theoretical frameworks heavily rely on the reification of mathematical models, patching together of disparate theoretical frames, tailoring some and axiomatising then in an opportunistic fashion. As we saw, EFTs heavily rely on such strategies and typically defy any residual intuitions of linearity in scientific progress. Thus, we may not know when exactly to stop digging⁵, but in the ebb and flow of top-down and bottom-up approaches to theory construction, we can't help but feel that currently, we probably need to cast our shovels aside and pick up our binoculars in search for fruitful insights across our horizon. Turning to the more destructive aspect of our discourse, we conclude that the skeptical remarks on naturalness and the resultant corroboration of the EFT framework on a casuistic basis underline the shortcoming of the project of developing scientific theories on the basis of some quasi transcendental principles. Of course, naturalness is only one instance of this. 20th century physics has seen multiple attempts (of limited success) to capture

⁵But see Crowther 2019 for a recent take on this question on the basis of some desiderata.

the physical content of a theory using a few fundamental principles or hypotheses. Examples, of course, include the two postulates of SR, the three principles (general covariance, general relativity and Mach's dictum) for GR, but also more metaphysical (in a quasi-transcendental sense) principles such as separability. Of course, a set of such a core of fundamental principles has frequently been suggested for QFT itself:

- Unitarity [from QM]
- Microcausality [from SR]
- Cluster Decomposition
- Renormalisability
- Naturalness

The revisability of some (e.g. renormalisability in light of EFTs or Lorentz invariance for curved spacetimes) or the potential abandonment of others (such as naturalness) has forcefully demonstrated that a quasi-rationalistic project of uncovering some basic principles pervading the description of the fundamental level of the universe is dispensable, overly ambitious or even misguided. This is where Einstein's realism probably went too far. In this sense, the flexibility-revisability of Quine's "web of belief" is vindicated. In the end, no single principle can safely be said to have cemented its status in the pantheon of physical principles. We have to abandon any aspiration for an "ante factum" aprioristic identification of the principles to be preserved in future frameworks.

At the same time, it is important to avoid over-stressing this point. For, despite their frequently pragmatic or heuristic birth, such principles do serve as the foundations for specific frameworks (recall our "hierarchical" structure of scientific theories in chapter 1.) that function as the lenses through which we view the world. The suggestion by Stein 2004 to take Carnap's frameworks along a more methodological path is pretty useful in this regard. One starts by constructing abstract, formalised-as-possible (mathematical) frames whose "essence" is frequently (but perhaps not uniquely – as the case of GR shows) captured by a set of key physical assumptions. These can often be elevated to the status of principles in the sense that their abandonment or (radical) revision implies a change in the domain a theory applies, the way it specifies physical systems, the way it treats dynamical variables variables etc. Describing a system in quantum mechanical terms differs from describing it in classical terms. Similarly, in relativistic theories

the precise of coordinates is often harder to come by than in classical mechanics. The dynamic expansion and continuous employment of a framework will help consolidate and cement the status of such principles, which often start out as “aporias” of a heuristic value.

From this perspective, we can appreciate how Lorentz invariance is a robust constitutive principle of the QFT framework while renormalisability was eventually modified in the light of EFTs. The status of “naturalness” is the most dubious of all; its systematic predictive failures have rendered the possibility of its complete abandonment pretty strong. By contrast, if the problematic around it were to be vindicated, we would be able to treat questions of naturalness in a realistic manner fully embracing their central role in theory construction. We can thus sympathise with physicists who, treating naturalness as a quasi-principle of the QFT framework, were frustrated to see that such elegant suggestions as SUSY, capable of resolving two major problems of modern physics with one stroke, have failed to garner any tangible evidence in their favour. Yet, from our point of view, this need only be construed as a negative example of the “dialectic” between framework principles and experience we have been expounding.

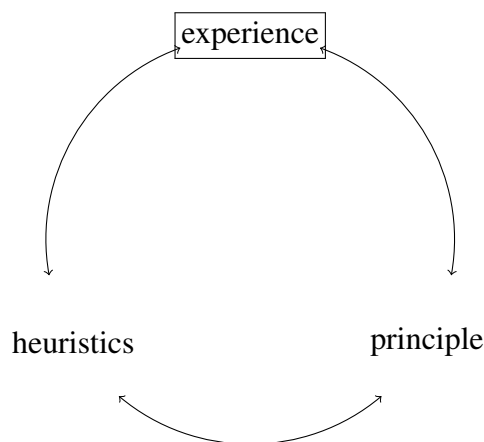


Figure 5.1: The Dialectic of Theory Construction

5.2.1 Theory Progression with the Help of EFTs

Within the broader meta-framework of EFTs, one treats theories as tentative, bound and approximate descriptions of a particular regime and a particular class of systems. Expanding and refining these descriptions is a procedure which, unless it fully breaks, the meta-framework can help systematise:

1. Specify the (class of) system(s) to be described: define the relevant degrees of freedom, the physical magnitudes that will play a role, the level of coarse-graining.

→*Example: Pick a system as simple as an electron interacting with a positron, write down the Lagrangian for the fermions, the photons exchanged, ignore other interactions such as gravity. In more complicated systems, like nuclei, one “forgets” about the internal structure of nucleons.*

2. Establish the energy scales of the problem to delimit, and estimate, the level of precision to be achieved.

→*Example: Integrate the electron out of the model to do non-relativistic interactions of low energy photons; or do not attempt to use QCD at energies low than some MeV*

3. Push the description to its limits; expand the description by including further correction terms and assess their behaviour: is the resulting model renormalisable? When does it break down? What new corrections should be taken into account?

→*Example: extensions of the standard model towards operators of dimension 5, 6 etc*

4. Depending on the outcome of the previous step, revise, abandon or adopt some principle to alter the descriptive framework.

→*Example: revise the scope of Lorentz invariance when extending QFT to curved spacetimes, abandon the vacuum in non-standard treatments of QFT in curved spacetimes, add new symmetries (from new gauge symmetries in SM to SUSY)*

Conclusion Revising a given theory or developing a new one is of course a pretty intricate business that cannot be done justice to when treated in a quasi-algorithmic procedure. Nonetheless, EFTs have allowed scientists to systematise the process of constructing new models on the basis of simple assumptions about the behaviour of physics across scales exploiting symmetries and their breaking. Yet, major breakthroughs will, in all likelihood, involve substantial revisions to the physical assumptions pervading the whole edifice of QFT. Along these lines, jettisoning naturalness might erode the ground on which aspiring BSM models such as SUSY are growing. Correspondingly, reassessing the status of demands such as UV completeness might help strengthen the QFT framework against successors such as string theory.

5.3 Methodology II: Extra-empirical Guiding Principles

In our discussion of naturalness, we saw how the problematic related to fine-tuning acts as a bridge between naturalness as autonomy and naturalness as typicality and, eventually, as a principle for theory selection. In this most recent “mutation”, naturalness more clearly showcases its role as an extra-empirical principle for comparing competing models in BSM physics. A lot of work in the area of supersymmetric, string-theoretic and other extensions or replacements of the SM have employed naturalness as a quantitative constraint on the likelihood of such models. This breaks some ground with our traditional understanding of confirming a physical model through (purely) experimental or observational means. In this sense, naturalness becomes part of a novel methodological standpoint that seeks to sidestep the need for (direct) empirical evidential support and reach a verdict about the cogency of a particular model using “peripheral” methodological assumptions.

I think that the situation bears a very close relation to a recent approach on the role of non-empirical confirmation for string theory [ST] developed by Dawid (Dawid 2013a, Dawid 2018). The core idea is very much the same: extend our methodological tools for theory confirmation into contexts where we lack experimental or observational data or where these are almost impossible to acquire. As is well known, ST is very problematic in precisely this respect since its main theoretical assumptions and predictions (such as additional spacetime dimensions, infinite vacuum states, supersymmetric particles) are beyond our experimental grasp: triggering the energy scales at which ST corrections to the SM play a role would require absurdly larger colliders than what is technologically physical in the immediate future (Conlon 2016, p. 107, Woit 2006, p. 208-16). Yet, certain researchers in the field not only continue their work unabashed, but also express full confidence in its eventual vindication. In fact, commitment to this research program is sometimes expressed in such an uncompromising manner that string theorists will not even back down from promoting a novel conception of science:

“The moment you encounter string theory and realize that almost all of the major developments in physics over the last hundred years emerge – and emerge with such elegance – from such a simple starting point, you realize that this incredibly compelling theory is in a class of its own”

—Green, quoted in Greene 2000 (p. 139)

This is often supported on the premise that string theory is a program unlike any other previ-

ously seen in physics in scope and ambition such that:

“There are no alternatives . . . all good ideas are part of string theory”

—Polchinski, quoted in Penrose 2016 (p. 87)

Dissenting voices are sometimes dismissed as echoes – relics of an old-fashioned and restrictive methodological standpoint, often in an admittedly harsh (yet undeniably amusing) manner:

“But the pontification, by the “Popperazzi”, about what is and is not science has become so furious in news reports and Internet blogs that I feel I have to address it [. . .] it would be the height of stupidity to dismiss a possibility just because it breaks some philosopher’s dictum about falsifiability”

—Susskind (in Susskind 2008, p. 192-4)

The dramatically novel character of string theory and the unmatched promise it holds as the unificatory project par excellence is, quite naturally, without peer in the too-slow-to-catch-on methodological discussions of philosophers. Therefore, the argument appears to be, to judge string theory by the appropriate standards, we must cast aside any methodological qualms inspired by old-fashioned philosophical prejudices such as falsifiability and adopt a methodology that will be up to the task of truly taking into account and evaluating the relevant merits of the grand program.

Dawid’s work (Dawid 2009, Dawid 2013a, Dawid 2013b, Dawid 2018) can be seen as a recent attempt to develop this line of thought in the philosophical literature and track its implications for theory confirmation. Dawid argues that the standards espoused by string theorists are aligned with a “meta-paradigmatic shift” in theory acceptance towards a new kind of confirmation theory, which he dubs “non-empirical confirmation” (NEC). NEC distances itself from the old ideal of assessing a theory’s status through direct confrontation with reality, emphasizing virtues other than empirical adequacy instead. Dawid (e.g. Dawid 2013b, p. 31-8) bases NEC on a triad of arguments or meta-principles which, despite being relatively weak on their own, can provide strong support for a theory when working in concert. According to this new approach, our degree of belief in a theory T is augmented whenever:

- **No Alternatives Argument (NAA):** despite sincere efforts, the scientific community has failed to produce alternatives to a theory T which successfully explains/solves a set of phenomena/problems that theory T is meant to encompass

- **Unexpected Explanatory Coherence Argument (UEA):** the theory manages to provide explanations and incorporate disparate phenomena and results within a coherent framework using a minimal set of postulates — theoretical posits
- **Meta-inductive Argument (MIA):** theories sharing the same principles as T have been confirmed; in other words, T is just using tools that have proven to be efficient in previously successful theories (i.e. T constitutes a continuation or natural extension of a successful paradigm)

Now, according to Dawid, our belief in ST is strengthened by each of the above arguments. For, in accordance with NAA, there is currently no alternative theory that aspires to provide a unified framework of all interactions — with contenders arising in the, admittedly, more limited project of quantising gravity (e.g. Kiefer 2007 2007, p. 279). In addition, ST was highly successful in bringing together various patches of high energy physics by incorporating supersymmetry, descriptions of black hole entropy and even “deriving” gravity as an interaction — proving that UEA holds. Finally, MIA is supposed to follow from the fact that ST can be viewed as an extension of the SM since it started out as an extension of the perturbative techniques that are very frequently employed in the SM context (e.g. Rickles 2016).

Dawid claims that arguments such as the above have been employed throughout the history of science to provide support for scientific theories. Theories like atomism, for instance, were dominant within the scientific community long before the advent of experiments that proved them right. But even in the case of experiments, Dawid continues, there is no way to decide upon whether they confirm the theory at hand, unless alternatives have been eliminated. Invoking work by Sklar 2000 and Stanford 2006 on scientific underdetermination (SU), Dawid suggests that all (historical) cases of theory confirmation are also cases of decision under underdetermination. This means that the scientists first need to significantly narrow the space of conceivable alternatives before they can be confident that the final verdict they reach (i.e. choosing one theory among competitors) is correct. It is here that NEC plays a vital role by limiting the potential rivals and guiding them to the right solution. One does not examine all possible models for the structure of atomic nuclei, for example, but rather chooses the ones that offer certain advantages such as simplifying calculations, conforming to what we know about the behaviour of protons and neutrons, the nature of interactions between the nucleus and electrons and so on.

So far so good. Here comes the provocative twist. Dawid goes on to suggest that, because

NEC has been playing a progressively more significant role in theory confirmation in High Energy Physics, there is no clear dichotomy between NEC and EC (the classical paradigm of theory confirmation). Rather, we can see theory confirmation as a continuous spectrum in which empirical testability is only one among many parameters that some theories might possess at a higher degree than others. ST can then be seen as occupying the most extreme end of the “theory confirmation spectrum” – the end at which trust in a theory is affirmed solely on the non-empirical criteria. Far from being unfalsifiable or speculative, string theory, from this perspective, simply illustrates in the most dramatic fashion the extremes of a certain style of reasoning that has been implicitly accepted in our theorising about the assessment of scientific theories since ever. The unique features of ST render it a theoretical frame without par and strongly suggest that, adjusting properly the weight we accord to extra-empirical virtues, we should take it seriously as a correct description of the world. In a nutshell, string theory is vindicated as the unique theoretical framework that emerges when we follow the logical progression of research in fundamental physics.

Assessment The instinctive reaction to the above arguments will probably be disbelief. One does not have to be a positivist or a Popperian to sense that things are moving pretty fast here. Fortunately, instinct can be put into words for at least two arguments can be presented against Dawid’s reasoning. In fact, both arguments may heavily draw on the insights and analyses of the EFT framework; the first is of a more historical flavour (essentially attacking MIA), while the other contests the (broader) assumption that meta-empirical assumptions might guarantee a (clear) verdict in favour of a theory.

As noted, an important component of Dawid’s argumentation is the affinity between the (scientific) fields of particle physics, as standardly practised within the QFT framework, and string theory. Indeed, string theory is often, and probably with good reason, taken to follow the tradition of particle physics – as evidenced by its adoption of similar assumptions (field-like structures evolving on a background structure⁶), same (mathematical) techniques (extending quantisation to strings and branes, using the RG to identify conformal theories, perturbation theory) and a similar final goal: the unification of all fundamental interactions. Yet, for all the conceptual and technical affinities between the two frameworks, it is also clear that important

⁶This is related to a point of contention between string theory and loop quantum gravity: the status of background independence. Proponents of the latter program have put emphasis on the fact that LQG continues GR’s abandonment of the need for some background spacetime structure – a feature to be preserved in any true successor theories.

dissimilarities exist between them, evident not least thanks to the effective nature of QFTs. Indeed, one cannot help but insist on the very fruitful interplay between theory and experiment that guided the whole development of the standard model:

- The development of regularisation and renormalisation techniques in the 1940s to save QED from its apparent demise. → prediction of the anomalous magnetic moment.
- The systematic employment of group theory to:
 - incorporate more interactions: e.g. expanding the model from electromagnetic to electroweak interactions – with the addition of the Higgs mechanism being an important piece to render the whole theory coherent. → masses of W^+ , W^- , Z^0 bosons.
 - describe autonomous degrees of freedom: mesons, hadrons etc. Central role of SU(3) symmetry → prediction of the Ω particle.
- Using renormalisability or its failure as a hint to developing new theories – e.g. the replacement of Fermi theory by electroweak theory. → seen to breakdown at energies higher than 100 GeV
- Treating QFTs as effective theories has allowed for more flexibility: one can now use a non-renormalisable theory to describe phenomena within a certain regime (even gravity!) and/or try to add new terms to existing ones. → adding neutrino masses to the SM, treating GR as a low-energy approximation to a more fundamental theory.

The advent of EFTs has made the tight connection between QFTs and experiment even more evident. Recall that bottom-up EFTs are constructed by first writing down a term expansion consistent with the symmetries and other constraints of the problem and then specifying the parameters using empirical data. This lets physicists work in ignorance of the more fundamental physical theory but in constant juxtaposition with experimental data. Now, string theories can (and do) engage in similar project of model construction, however these can at least be assessed on the basis of extra-empirical principles like naturalness (like SUSY models) and only once the low-energy limit can be established – so that they act as corrections to known results. Despite some conceptual and mathematical affinities, it is far-fetched to claim that QFT and ST share enough common ground. At least from a more historical and methodological standpoint, this is simply inaccurate.

Now, a supporter of MIA could simply insist that the only continuity we need here is that of the (mathematical and conceptual) tools employed in theory T and its successor, so to insist on their diverging epistemic standing simply amounts to no more than begging the question. There might be something to this, but I think that apart from the actual difference in practice, the whole reasoning of MIA, in this context at least, is highly problematic. The main concern with extra-empirical principles for theory choice and construction, as we saw when examining naturalness, is that they are liable to being tweaked so as to allow the kind of verdict one desires. This is, for example, the problem with rather vague criteria such as simplicity or “elegance”, especially when empirical evidence is unavailable. The case of SUSY is telling. One could not have asked for a better or elegant solution to problems such as the hierarchy or cosmological constant. Yet, for all the serendipity, evidence in favour of the theory has so far eluded us. One can see SUSY as a beta form of string theory: the theory provided unified solutions explanations to a series of disparate problems, continued on the path of unification most fully employing the techniques of particle physics and even lacked any competitor (with the exception of string theory, which is after all meant to reproduce it as a low energy limit). In other words, in SUSY we have the most dramatic instance of thwarted expectations according to NEC. To paraphrase Kant with a twist one can claim that “empirical results without theoretical principles are often blind but theoretical principles without empirical results often become blinding”.

Conclusion The SM of particle physics has been one of, if not the, most successful theory ever devised by man. Its consistent and spectacular predictive success has been the culmination of a century-long exercise in careful model building. Throughout this period, sophisticated and insightful theoretical analysis found its match in systematic and precise experimentation. Frequently starting from phenomenological models, physicists were able to divine the right fundamental laws at work combining insights about the symmetries at play, the behaviour of fields at different scales and the like. QFT itself, throughout the course of its history, often faced the bleak prospect of abandonment – only to be saved at the last minute by brilliant breakthroughs such as those in 40s (renormalisation), 60s (symmetry breaking) and 70s (renormalisation group). Slow and careful tactics with some masterstrokes of ingenuity were the source of progress throughout the 20th century research in fundamental physics. Consequently, when our intuition about the “inner workings” of nature is lacking or cannot be trusted, it is perhaps more prudent to postpone judgment about the truth of a theory up until further evidence can be

gathered.

References

- 't Hooft, Gerard (1980). “Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking”. In: *NATO Sci. Ser. B* 59. Ed. by Gerard 't Hooft et al., pp. 135–157. DOI: 10.1007/978-1-4684-7571-5_9.
- Adlam, Emily (2019). “Fundamental?” In: *What is Fundamental?* Ed. by Anthony Aguirre, Brendan Foster, and Zeeya Merali. Springer Verlag, pp. 5–12.
- Agon, Cesar et al. (2018). “Coarse Grained Quantum Dynamics”. In: *Phys. Rev. D* 98.2, p. 025019. DOI: 10.1103/PhysRevD.98.025019. arXiv: 1412.3148 [hep-th].
- Aitchison, Ian J. R. and Anthony J. G. Hey (2013). *Gauge Theories in Particle Physics: A Practical Introduction, Volume 1 : From Relativistic Quantum Mechanics to QED, Fourth Edition*. Taylor & Francis. ISBN: 978-1-4665-1302-0, 978-0-429-18538-0, 978-1-4665-1299-3. DOI: 10.1201/b13717.
- Albert, David Z. (1996). “Elementary Quantum Metaphysics”. In: *Bohmian Mechanics and Quantum theory: An Appraisal*. Ed. by J. T. Cushing, Arthur Fine, and Sheldon Goldstein. Kluwer Academic Publishers, pp. 277–284.
- (1992). *Quantum Mechanics and Experience*. Harvard Up.
- Allori, Valia (2013). “Primitive Ontology and the Structure of Fundamental Physical Theories”. In: *The Wave Function: Essays in the Metaphysics of Quantum Mechanics*. Ed. by Alyssa Ney and David Z. Albert. Oxford University Press.
- Alvarez-Gaumé, L. and M.A. Vázquez-Mozo (2011). *An Invitation to Quantum Field Theory*. Lecture Notes in Physics. Springer Berlin Heidelberg. ISBN: 9783642237270. URL: <https://books.google.gr/books?id=0oNj1AdbFq8C>.
- Amelino-Camelia, Giovanni et al. (Oct. 2011). “Principle of relative locality”. In: *Phys. Rev. D* 84 (8), p. 084010. DOI: 10.1103/PhysRevD.84.084010. URL: <https://link.aps.org/doi/10.1103/PhysRevD.84.084010>.

- Anderson, Greg W. and Diego J. Castano (1995). “Measures of fine tuning”. In: *Phys. Lett. B* 347, pp. 300–308. DOI: 10.1016/0370-2693(95)00051-L. arXiv: hep-ph/9409419.
- Appelquist, Thomas and J. Carazzone (May 1975). “Infrared singularities and massive fields”. In: *Phys. Rev. D* 11 (10), pp. 2856–2861. DOI: 10.1103/PhysRevD.11.2856. URL: <https://link.aps.org/doi/10.1103/PhysRevD.11.2856>.
- Bain, Jonathan (2013). “Effective Field Theories”. In: *The Oxford Handbook of Philosophy of Physics*. Ed. by Robert W. Batterman. Oup Usa, p. 224.
- (2019). “Why Be Natural?” In: *Foundations of Physics* 49.9, pp. 898–914. DOI: 10.1007/s10701-019-00249-z.
- Baker, Mark Robert, Natalia Kiriushcheva, and Sergei V. Kuzmin (2021). “Noether and Hilbert (metric) energy-momentum tensors are not, in general, equivalent”. In: *Nuclear Physics* 962, p. 115240.
- Barbieri, R. and Gian Francesco Giudice (1988). “Upper bounds on supersymmetric particle masses”. In: *Nuclear Physics B* 306.1, pp. 63–76. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(88\)90171-X](https://doi.org/10.1016/0550-3213(88)90171-X). URL: <https://www.sciencedirect.com/science/article/pii/055032138890171X>.
- Batterman, Robert W. (2017). “Philosophical Implications of Kadanoff’s Work on the Renormalization Group”. In: *Journal of Statistical Physics* 167.3-4, pp. 559–574. DOI: 10.1007/s10955-016-1659-9.
- (2001). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford University Press.
- (2013). “The Tyranny of Scales”. In: *The Oxford handbook of philosophy of physics*. Ed. by Robert W. Batterman. Oxford University Press, pp. 255–286.
- Becher, T. (2015). “Modern Effective Field Theory, Lecture Notes”. In: URL: http://freeuni.edu.ge/physics/physics_1/index_files/files/modernEFT_I.pdf.
- Belot, Gordon (2007). “Is Classical Electrodynamics an Inconsistent Theory?” In: *Canadian Journal of Philosophy* 37.2, pp. 263–282. DOI: 10.1353/cjp.2007.0014.
- Bianchi, Eugenio and Carlo Rovelli (Feb. 2010). “Why all these prejudices against a constant?” In: arXiv: 1002.3966 [astro-ph.CO].

- Birrell, N.D. and P.C.W. Davies (1984). *Quantum Fields in Curved Space*. Cambridge Monographs on Mathematical Physics. Cambridge University Press. ISBN: 9780521278584. URL: <https://books.google.gr/books?id=SEnaUnrqzrUC>.
- Bishop, Robert (2017). “Chaos”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2017. Metaphysics Research Lab, Stanford University.
- Bohr, Niels and Leon Rosenfeld (1933). “On the Question of the Measurability of Electromagnetic Field Quantities”. In: DOI: 10.1007/978-94-009-9349-5_26.
- (1996). “V - Field and Charge Measurements in Quantum Electrodynamics”. In: *Foundations of Quantum Physics II (1933–1958)*. Ed. by Jørgen Kalckar(ed). Vol. 7. Niels Bohr Collected Works. Elsevier, pp. 211–216. DOI: [https://doi.org/10.1016/S1876-0503\(08\)70368-2](https://doi.org/10.1016/S1876-0503(08)70368-2). URL: <https://www.sciencedirect.com/science/article/pii/S1876050308703682>.
- Bokulich, Alisa (2008). *Reexamining the Quantum-Classical Relation: Beyond Reductionism and Pluralism*. Cambridge University Press.
- Borrelli, Arianna and Elena Castellani (2019). “The Practice of Naturalness: A Historical-Philosophical Perspective”. In: *Foundations of Physics* 49.9, pp. 860–878. DOI: 10.1007/s10701-019-00287-7.
- Boscovich, R.G. and J.M. Child (2018). *A Theory of Natural Philosophy*. Creative Media Partners, LLC. ISBN: 9780342853397. URL: <https://books.google.gr/books?id=mXhDvgEACAAJ>.
- Brahma, Suddhasattwa, Omar Alaryani, and Robert Brandenberger (Aug. 2020). “Entanglement entropy of cosmological perturbations”. In: *Phys. Rev. D* 102 (4), p. 043529. DOI: 10.1103/PhysRevD.102.043529. URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.043529>.
- Brandenberger, Robert (Aug. 2021). “Limitations of an Effective Field Theory Treatment of Early Universe Cosmology”. In.
- Brown, L.M. (1994). *Renormalization: From Lorentz to Landau (and Beyond)*. Springer London, Limited. ISBN: 9783540944010. URL: <https://books.google.gr/books?id=q4UIGgAACAAJ>.
- Burgess, Cliff P. (Nov. 2007). “An Introduction to Effective Field Theory”. In: *Annual Review of Nuclear and Particle Science* 57.1, pp. 329–362. ISSN: 1545-4134. DOI: 10.1146/

annurev.nucl.56.080805.140508. URL: <http://dx.doi.org/10.1146/annurev.nucl.56.080805.140508>.

- Burgess, Cliff P. (2020). *Introduction to Effective Field Theory: Thinking Effectively about Hierarchies of Scale*. Cambridge University Press. DOI: 10.1017/9781139048040.
- (2004). “Quantum Gravity in Everyday Life: General Relativity as an Effective Field Theory”. In: *Living Reviews in Relativity* 7.
- Burgess, Cliff P. and Guy D. Moore (2006). *The Standard Model: A Primer*. Cambridge University Press. ISBN: 9780511256523. URL: <https://books.google.gr/books?id=o32xAQAACAAJ>.
- Butterfield, Jeremy (2011). “Less is Different: Emergence and Reduction Reconciled”. In: *Foundations of Physics* 41.6, pp. 1065–1135. DOI: 10.1007/s10701-010-9516-1.
- Cao, Tian Yu (1999). *Conceptual Foundations of Quantum Field Theory*. Cambridge University Press.
- Cao, Tian Yu and Silvan S. Schweber (1993). “The Conceptual Foundations and the Philosophical Aspects of Renormalization Theory”. In: *Synthese* 97.1, pp. 33–108. DOI: 10.1007/BF01255832.
- Carnap, Rudolf (1974). *An Introduction to the Philosophy of Science*. Dover Publications.
- Carroll, S.M. (2019). *Spacetime and Geometry*. Cambridge University Press. ISBN: 9781108488396. URL: <https://books.google.gr/books?id=PTGdDwAAQBAJ>.
- Cartwright, Nancy (1983). *How the Laws of Physics Lie*. Oxford, England: Oxford University Press.
- (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press.
- Castellani, Elena (June 2002). “Reductionism, Emergence, and Effective Field Theories”. In: *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics* 33, pp. 251–267. DOI: 10.1016/S1355-2198(02)00003-5.
- Chakravartty, Anjan (2017). “Scientific Realism”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2017. Metaphysics Research Lab, Stanford University.
- Clifton, Rob and Hans Halvorson (2001). “Are Rindler Quanta Real? Inequivalent Particle Concepts in Quantum Field Theory”. In: *British Journal for the Philosophy of Science* 52.3, pp. 417–470. DOI: 10.1093/bjps/52.3.417.

- Conlon, J. (2016). *Why String Theory?* CRC Press. ISBN: 9781482242492. URL: <https://books.google.gr/books?id=F5DwCgAAQBAJ>.
- Craig, Nathaniel (2017). “Naturalness and New Approaches to the Hierarchy Problem”. In: URL: <https://www.ias.edu/sites/default/files/pitp/craig.pdf>.
- Craig, Nathaniel and Seth Koren (2020). “IR Dynamics from UV Divergences: UV/IR Mixing, NCFT, and the Hierarchy Problem”. In: *JHEP* 03, p. 037. DOI: 10.1007/JHEP03(2020)037. arXiv: 1909.01365 [hep-ph].
- Crowther, Karen (2016). *Effective Spacetime: Understanding Emergence in Effective Field Theory and Quantum Gravity*. Springer.
- (2019). “When Do We Stop Digging? Conditions on a Fundamental Theory of Physics”. In: *What is ‘Fundamental’?* Ed. by Anthony Aguirre, Brendan Foster, and Zeeya Merali. Springer, pp. 123–133.
- Crowther, Karen and Niels Linnemann (2019). “Renormalizability, Fundamentality, and a Final Theory: The Role of UV-Completion in the Search for Quantum Gravity”. In: *British Journal for the Philosophy of Science* 70.2, pp. 377–406. DOI: 10.1093/bjps/axx052.
- Cugnon, J. (2012). “The Casimir effect and the vacuum energy: Duality in the physical interpretation”. In: *Few Body Syst.* 53. Ed. by J. R. Cudell, J. P. Lansberg, and D. Mancusi, pp. 181–188. DOI: 10.1007/s00601-011-0250-9.
- Curiel, Erik (2017). “A Primer on Energy Conditions”. In: *Towards a Theory of Spacetime Theories*. Ed. by Dennis Lehmkuhl, Gregor Schiemann, and Erhard Scholz. Birkhäuser, pp. 43–104.
- (2014). “Classical Mechanics Is Lagrangian; It Is Not Hamiltonian”. In: *British Journal for the Philosophy of Science* 65.2, pp. 269–321. DOI: 10.1093/bjps/axs034.
- (Mar. 2019). “Schematizing the Observer and the Epistemic Content of Theories”. In: Dawid, Richard (2018). “Delimiting the Unconceived”. In: *Foundations of Physics* 48.5, pp. 492–506. DOI: 10.1007/s10701-017-0132-1.
- (2009). “On the Conflicting Assessments of the Current Status of String Theory”. In: *Philosophy of Science* 76.5, pp. 984–996. DOI: 10.1086/605794.
- (2013a). *String Theory and the Scientific Method*. Cambridge University Press.
- (2013b). “Theory Assessment and Final Theory Claim in String Theory”. In: *Foundations of Physics* 43.1, pp. 81–100. DOI: 10.1007/s10701-011-9592-x.
- Deutsch, David (1997). *The Fabric of Reality*. Allan Lane.

- Dicke, R. H. (Nov. 1961). “Dirac’s Cosmology and Mach’s Principle”. In: *Nature* 192.3512, pp. 440–441. DOI: 10.1038/192440a0.
- Dine, Michael (Jan. 2015). “Naturalness Under Stress”. In: *Annual Review of Nuclear and Particle Science* 65. DOI: 10.1146/annurev-nucl-102014-022053.
- Dirac, Paul. A. M. (Feb. 1937). “The Cosmological Constants”. In: 139.3512, p. 323. DOI: 10.1038/139323a0.
- Disalle, Robert (1995). “Spacetime Theory as Physical Geometry”. In: *Erkenntnis* 42.3, pp. 317–337. DOI: 10.1007/bf01129008.
- Donoghue, John (2017). “Quantum gravity as a low energy effective field theory”. In: *Scholarpedia* 12.4, p. 32997. DOI: 10.4249/scholarpedia.32997.
- Duncan, A. (2012). *The Conceptual Framework of Quantum Field Theory*. EBSCO ebook academic collection. OUP Oxford. ISBN: 9780199573264. URL: <https://books.google.gr/books?id=Muh0TQvpY5sC>.
- Dürr, D. and S. Teufel (2009). *Bohmian Mechanics: The Physics and Mathematics of Quantum Theory*. Springer Berlin Heidelberg. ISBN: 9783540893448. URL: <https://books.google.gr/books?id=UWP2ZSs-UD0C>.
- Dyson, F. J. (Feb. 1952). “Divergence of Perturbation Theory in Quantum Electrodynamics”. In: *Phys. Rev.* 85 (4), pp. 631–632. DOI: 10.1103/PhysRev.85.631. URL: <https://link.aps.org/doi/10.1103/PhysRev.85.631>.
- Earman, John (2004). “Laws, Symmetry, and Symmetry Breaking: Invariance, Conservation Principles, and Objectivity”. In: *Philosophy of Science* 71.5, pp. 1227–1241. DOI: 10.1086/428016.
- Earman, John and Doreen Fraser (2006). “Haag’s Theorem and its Implications for the Foundations of Quantum Field Theory”. In: *Erkenntnis* 64.3, pp. 305–344. DOI: 10.1007/s10670-005-5814-y.
- Ellis, George (Feb. 2012). “Top-down causation and emergence: Some comments on mechanisms”. In: *Interface focus* 2, pp. 126–40. DOI: 10.1098/rsfs.2011.0062.
- Feng, Jonathan L. (2013). “Naturalness and the Status of Supersymmetry”. In: *Annual Review of Nuclear and Particle Science* 63, pp. 351–382.
- Feyerabend, Paul K. (1962). “Explanation, Reduction and Empiricism”. In: *Crítica: Revista Hispanoamericana de Filosofía*. Ed. by H. Feigl and G. Maxwell, pp. 103–106.

- Feynman R., P. (2014). *QED: The Strange Theory of Light and Matter*. Princeton Science Library. Princeton University Press. ISBN: 9781400847464. URL: <https://books.google.com/books?id=2o2JfTDiA40C>.
- Fichet, S. (Dec. 2012). “Quantified naturalness from Bayesian statistics”. In: *Phys. Rev. D* 86 (12), p. 125029. DOI: 10.1103/PhysRevD.86.125029. URL: <https://link.aps.org/doi/10.1103/PhysRevD.86.125029>.
- Fisher, Michael E. (Apr. 1998). “Renormalization group theory: Its basis and formulation in statistical physics”. In: *Rev. Mod. Phys.* 70 (2), pp. 653–681. DOI: 10.1103/RevModPhys.70.653. URL: <https://link.aps.org/doi/10.1103/RevModPhys.70.653>.
- Folland, G.B. (2013). *Quantum Field Theory: A Tourist Guide for Mathematicians*. Mathematical surveys and monographs. American Mathematical Society. URL: <https://books.google.gr/books?id=KWMAoQEACAAJ>.
- Fradkin, Eduardo (2021). *Quantum Field Theory: An Integrated Approach*. Princeton University Press. ISBN: 9780691189550. URL: <https://books.google.gr/books?id=quEIEAAAQBAJ>.
- Franklin, Alexander (2020). “Whence the Effectiveness of Effective Field Theories?” In: *British Journal for the Philosophy of Science* 71.4, pp. 1235–1259. DOI: 10.1093/bjps/axy050.
- Fraser, Doreen (2011). “How to Take Particle Physics Seriously: A Further Defence of Axiomatic Quantum Field Theory”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 42.2, pp. 126–135. DOI: 10.1016/j.shpsb.2011.02.002.
- (2009). “Quantum Field Theory: Underdetermination, Inconsistency, and Idealization”. In: *Philosophy of Science* 76.4, pp. 536–567. DOI: 10.1086/649999.
- (2012). “Spontaneous Symmetry Breaking: Quantum Statistical Mechanics Versus Quantum Field Theory”. In: *Philosophy of Science* 79.5, pp. 905–916. DOI: 10.1086/668044.
- (2020). “The Development of Renormalization Group Methods for Particle Physics: Formal Analogies Between Classical Statistical Mechanics and Quantum Field Theory”. In: *Synthese* 197.7, pp. 3027–3063. DOI: 10.1007/s11229-018-1862-0.

- Fraser, Doreen (2008). “The Fate of ‘Particles’ in Quantum Field Theories with Interactions”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 39.4, pp. 841–859. DOI: 10.1016/j.shpsb.2008.05.003.
- Fraser, Doreen and Adam Koberinski (2016). “The Higgs Mechanism and Superconductivity: A Case Study of Formal Analogies”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 55, pp. 72–91. DOI: 10.1016/j.shpsb.2016.08.003.
- Fraser, James D. (2020a). “The Real Problem with Perturbative Quantum Field Theory”. In: *British Journal for the Philosophy of Science* 71.2, pp. 391–413. DOI: 10.1093/bjps/axx042.
- (2021). “The Twin Origins of Renormalization Group Concepts”. In: *Studies in History and Philosophy of Science Part A* 89, pp. 114–128. DOI: 10.1016/j.shpsa.2021.08.002.
- (Feb. 2020b). “Towards a Realist View of Quantum Field Theory”. In: *Scientific Realism and the Quantum*. Oxford University Press, pp. 276–292. DOI: 10.1093/oso/9780198814979.003.0014. URL: <https://oxford.universitypressscholarship.com/10.1093/oso/9780198814979.001.0001/oso-9780198814979-chapter-14>.
- (July 2016). “What is Quantum Field Theory? Idealisation, Explanation and Realism in High Energy Physics.” URL: <https://etheses.whiterose.ac.uk/14415/>.
- (Dec. 2018). “Renormalization and the formulation of scientific realism.” In: *Philosophy of science*. 85.5, pp. 1164–1175. URL: <http://dro.dur.ac.uk/27619/>.
- Fredenhagen, Klaus, Karl-Henning Rehren, and Erhard Seiler (2007). “Quantum field theory: Where we are”. In: *Lect. Notes Phys.* 721, pp. 61–87. DOI: 10.1007/978-3-540-71117-9_4. arXiv: hep-th/0603155.
- French, Steven and D’Ecio Krause (2006). *Identity in Physics: A Historical, Philosophical, and Formal Analysis*. Oxford University Press.
- Friedman, Michael (1974). “Explanation and Scientific Understanding”. In: *Journal of Philosophy* 71.1, pp. 5–19. DOI: 10.2307/2024924.
- Frigg, Roman (2020). *Modelling nature : an opinionated introduction to scientific representation*. eng. Synthese library, v. 427. Cham, Switzerland: Springer. ISBN: 3030451534.

- Frigg, Roman and James Nguyen (Jan. 2020). “Modelling Nature: An Opinionated Introduction to Scientific Representation”. In: DOI: 10.1007/978-3-030-45153-0.
- Frisch, Mathias (2005). *Inconsistency, Asymmetry and Non Locality: A Philosophical Investigation of Classical Electrodynamics*. Oxford University Press.
- Fuchs, Christopher (Dec. 2017). “On Participatory Realism”. In: pp. 113–134. ISBN: 978-3-319-43758-3. DOI: 10.1007/978-3-319-43760-6_7.
- Georgi, H. (1993). “Effective field theory”. In: *Ann. Rev. Nucl. Part. Sci.* 43, pp. 209–252. DOI: 10.1146/annurev.ns.43.120193.001233.
- Ghirardi, G. C., R. Grassi, and F. Benatti (1995). “Describing the Macroscopic World: Closing the Circle Within the Dynamical Reduction Program”. In: *Foundations of Physics* 25.1, pp. 5–38. DOI: 10.1007/BF02054655.
- Giere, Ronald N. (2004). “How Models Are Used to Represent Reality”. In: *Philosophy of Science* 71.5, pp. 742–752. DOI: 10.1086/425063.
- (2006). *Scientific Perspectivism*. University of Chicago Press.
- (1999). “Using Models to Represent Reality”. In: *Model-Based Reasoning in Scientific Discovery*. Ed. by L. Magnani, N. J. Nersessian, and P. Thagard. Kluwer/Plenum, pp. 41–57.
- Giudice, Gian Francesco (June 2008). “Naturally Speaking: The Naturalness Criterion and Physics at the LHC”. In: DOI: 10.1142/9789812779762_0010.
- (2019). “The Dawn of the Post-Naturalness Era”. In: *From My Vast Repertoire ...: Guido Altarelli’s Legacy*. Ed. by Aharon Levy, Stefano Forte, and Giovanni Ridolfi, pp. 267–292. DOI: 10.1142/9789813238053_0013. arXiv: 1710.07663 [physics.hist-ph].
- Glashow, S. L., J. Iliopoulos, and L. Maiani (Oct. 1970). “Weak Interactions with Lepton-Hadron Symmetry”. In: *Phys. Rev. D* 2 (7), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285. URL: <https://link.aps.org/doi/10.1103/PhysRevD.2.1285>.
- Goldstein, H., C.P. Poole, and J.L. Safko (2002). *Classical Mechanics*. Addison Wesley. ISBN: 9780201657029. URL: <https://books.google.gr/books?id=tJCuQgAACAAJ>.
- Graña, Mariana and Alvaro Herráez (2021). *The Swampland Conjectures: A bridge from Quantum Gravity to Particle Physics*. arXiv: 2107.00087 [hep-th].
- Greene, B. (2000). *The Elegant Universe: Superstrings, Hidden Dimensions, and the Quest for the Ultimate Theory*. Vintage Series. Vintage Books. ISBN: 9780375708114. URL: <https://books.google.gr/books?id=8QYs7n3klvYC>.

- Grimmer, Daniel, Bruno de S. L. Torres, and Eduardo Martin-Martinez (Oct. 2021). “Measurements in QFT: Weakly coupled local particle detectors and entanglement harvesting”. In: *Phys. Rev. D* 104 (8), p. 085014. DOI: 10.1103/PhysRevD.104.085014. URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.085014>.
- Grinbaum, Alexei (2012). “Which Fine-Tuning Arguments Are Fine?” In: *Foundations of Physics* 42.5, pp. 615–631. DOI: 10.1007/s10701-012-9629-9.
- Gründler, Gerold Dipl.-Phys. (2013). “The Casimir-Effect: No Manifestation of Zero-Point Energy”. In: *arXiv: General Physics*.
- Guth, Alan H (June 2007). “Eternal inflation and its implications”. In: *Journal of Physics A: Mathematical and Theoretical* 40.25, pp. 6811–6826. DOI: 10.1088/1751-8113/40/25/s25. URL: <https://doi.org/10.1088/1751-8113/40/25/s25>.
- Hagar, Amit (2014). *Discrete or Continuous? The Quest for Fundamental Length in Modern Physics*. Cambridge University Press.
- Halvorson, Hans (2019). “To Be a Realist About Quantum Theory”. In: *Quantum Worlds: Perspectives on the Ontology of Quantum Mechanics*. Ed. by Olimpia Lombardi.
- Halvorson, Hans and Michael Mueger (2006). “Algebraic Quantum Field Theory”. In: *Handbook of Philosophy of Science: Philosophy of Physics*. Ed. by J. Butterfield and J. Earman. Kluwer Academic Publishers.
- Hancox-Li, Leif (2017). “Solutions in Constructive Field Theory”. In: *Philosophy of Science* 84.2, pp. 335–358. DOI: 10.1086/690722.
- Harlander, Robert and Joshua Rosaler (2019). “Higgs Naturalness and Renormalized Parameters”. In: *Foundations of Physics* 49.9, pp. 879–897. DOI: 10.1007/s10701-019-00296-6.
- Hartmann, Stephan (2001). “Effective Field Theories, Reductionism and Scientific Explanation”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 32.2, pp. 267–304. DOI: 10.1016/s1355-2198(01)00005-3.
- Healey, Richard (2017). *The Quantum Revolution in Philosophy*. Oxford University Press.
- Hebecker, Arthur (Jan. 2021). *Naturalness, String Landscape and Multiverse: A Modern Introduction with Exercises*. ISBN: 978-3-030-65150-3. DOI: 10.1007/978-3-030-65151-0.

- Hempel, Carl Gustav (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: The Free Press.
- Hollands, Stefan and Robert M. Wald (2008). “Quantum field theory in curved spacetime, the operator product expansion, and dark energy”. In: *Gen. Rel. Grav.* 40, pp. 2051–2059. DOI: 10.1007/s10714-008-0672-y. arXiv: 0805.3419 [gr-qc].
- (2015). “Quantum fields in curved spacetime”. In: *Phys. Rept.* 574, pp. 1–35. DOI: 10.1016/j.physrep.2015.02.001. arXiv: 1401.2026 [gr-qc].
- Hossenfelder, Sabine (2018). *Lost in Math: How Beauty Leads Physics Astray*. Basic Books. ISBN: 9780465094264. URL: <https://books.google.gr/books?id=0lc4DwAAQBAJ>.
- (2019). “Screams for Explanation: Finetuning and Naturalness in the Foundations of Physics”. In: *Synthese* 198.Suppl 16, pp. 3727–3745. DOI: 10.1007/s11229-019-02377-5.
- Howard, Don A. and Marco Giovanelli (2019). “Einstein’s Philosophy of Science”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2019. Metaphysics Research Lab, Stanford University.
- Huggett, Nick and Robert Weingard (1995). “The Renormalisation Group and Effective Field Theories”. In: *Synthese* 102.1, pp. 171–194. DOI: 10.1007/BF01063904.
- Jaffe, R. (Mar. 2005). “The Casimir Effect and the Quantum Vacuum”. In: *Physical Review D* 72. DOI: 10.1103/PhysRevD.72.021301.
- Jegerlehner, F. (Mar. 2015). “The hierarchy problem and the cosmological constant problem in the Standard Model”. In.
- Kiefer, C. (2007). *Quantum Gravity*. International Series of Monographs on Physics. OUP Oxford. ISBN: 9780199212521. URL: <https://books.google.gr/books?id=zLDRIqgjrw4C>.
- Kitcher, Philip (1989). “Explanatory Unification and the Causal Structure of the World”. In: *Scientific Explanation*. Ed. by Philip Kitcher and Wesley Salmon. Minneapolis: University of Minnesota Press, pp. 410–505.
- Koberinski, Adam (forthcoming). “Problems with the Cosmological Constant Problem”. In: *Philosophy Beyond Spacetime*. Ed. by Christian Wüthrich, Baptiste Le Bihan, and Nick Huggett. Oxford University Press.
- (2021). “Regularizing (Away) Vacuum Energy”. In: *Foundations of Physics* 51.1, pp. 1–22. DOI: 10.1007/s10701-021-00442-z.

- Koberinski, Adam and Christopher Smeenk (forthcoming). “Rethinking the cosmological constant problem”. In.
- Koren, Seth (2020). “New Approaches to the Hierarchy Problem and their Signatures from Microscopic to Cosmic Scales”. PhD thesis. UC, Santa Barbara (main).
- Kragh, H. (1990). *Dirac: A Scientific Biography*. Cambridge University Press. ISBN: 9780521380898. URL: <https://books.google.gr/books?id=zXm1Bs01VREC>.
- (2011). *Higher Speculations: Grand Theories and Failed Revolutions in Physics and Cosmology*. OUP Oxford. ISBN: 9780191003349. URL: <https://books.google.gr/books?id=3CSQDwAAQBAJ>.
- Krause, Decio and Jonas R. B. Arenhart (2016). *The Logical Foundations of Scientific Theories. Languages, Structures, and Models*. Nova Iorque, NY, EUA: Routledge.
- Kuhn, Thomas S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.
- Lancaster, Tom and Stephen J Blundell (Apr. 2014). *Quantum field theory for the gifted amateur*. Oxford: Oxford University Press. DOI: 0199699321. URL: <https://cds.cern.ch/record/1629337>.
- Landsman, K. (2017). *Foundations of Quantum Theory: From Classical Concepts to Operator Algebras*. Fundamental Theories of Physics. Springer International Publishing. ISBN: 9783319517773. URL: <https://books.google.gr/books?id=INpCDwAAQBAJ>.
- Laudan, Larry (1981). “A Confutation of Convergent Realism”. In: *Philosophy of Science* 48.1, pp. 19–49. DOI: 10.1086/288975.
- Lewis, David (1994). “Humean Supervenience Debugged”. In: *Mind* 103.412, pp. 473–490. DOI: 10.1093/mind/103.412.473.
- Li, Bihui (2015). “Coarse-Graining as a Route to Microscopic Physics: The Renormalization Group in Quantum Field Theory”. In: *Philosophy of Science* 82.5, pp. 1211–1223. DOI: 10.1086/684085.
- Loewer, Barry (1996). “Humean Supervenience”. In: *Philosophical Topics* 24.1, pp. 101–127. DOI: 10.5840/philtopics199624112.
- Malament, David (1996). “In Defense of Dogma: Why There Cannot Be a Relativistic Quantum Mechanical Theory of (Localizable) Particles”. In: *Perspectives on Quantum Reality*. Ed. by R. Clifton. Kluwer Academic Publishers.

- Manohar, Aneesh V. (Apr. 2018). “Introduction to Effective Field Theories”. In: ed. by Sacha Davidson et al. DOI: 10.1093/oso/9780198855743.003.0002. arXiv: 1804.05863 [hep-ph].
- Martin, Jérôme (July 2012). “Everything you always wanted to know about the cosmological constant problem (but were afraid to ask)”. In: *Comptes Rendus Physique* 13.6-7, pp. 566–665. ISSN: 1631-0705. DOI: 10.1016/j.crhy.2012.04.008. URL: <http://dx.doi.org/10.1016/j.crhy.2012.04.008>.
- Martin-Martinez, Eduardo (Nov. 2015). “Causality issues of particle detector models in QFT and quantum optics”. In: *Phys. Rev. D* 92 (10), p. 104019. DOI: 10.1103/PhysRevD.92.104019. URL: <https://link.aps.org/doi/10.1103/PhysRevD.92.104019>.
- Maudlin, Tim (Mar. 2007). “Completeness, supervenience and ontology”. In: *Journal of Physics A: Mathematical and Theoretical* 40, p. 3151. DOI: 10.1088/1751-8113/40/12/S16.
- (2019). *Philosophy of Physics: Quantum Theory*. Princeton University Press.
- Miller, Michael (Dec. 2016). *Mathematical Structure and Empirical Content*. URL: <http://philsci-archive.pitt.edu/12678/>.
- Morrison, Margaret (2015). *Reconstructing Reality: Models, Mathematics, and Simulations*. Oup Usa.
- Mosterin, Jesus (2004). “Anthropic Explanations in Cosmology”.
- Mumford, Stephen (2002). *Laws in Nature*. Routledge.
- Myrvold, Wayne C. (2021). *Beyond Chance and Credence: A Theory of Hybrid Probabilities*. Oxford University Press. ISBN: 9780198865094. URL: <https://books.google.com/books?id=6PruzQEACAAJ>.
- (2019a). “How Could Relativity Be Anything Other Than Physical”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 67, pp. 137–143. DOI: 10.1016/j.shpsb.2017.05.007.
- (2019b). “Ontology for Relativistic Collapse Theories”. In: *Quantum Worlds*.
- Newton, I. (2014). *The Principia: Mathematical Principles of Natural Philosophy*. University of California Press. ISBN: 9780520960916. URL: https://books.google.com/books?id=2u%5C_DBAAAQBAJ.
- Ney, Alyssa (2019). “Locality and Wave Function Realism”. In: *Quantum Worlds*.

- Nikolić, Hrvoje (2017). “Is zero-point energy physical? A toy model for Casimir-like effect”. In: *Annals of Physics* 383, pp. 181–195.
- (2016). “Proof that Casimir force does not originate from vacuum energy”. In: *Physics Letters B* 761, pp. 197–202.
- North, Jill (2007). “Mathias Frisch, Inconsistency, Asymmetry, and Non-Locality: A Philosophical Investigation of Classical Electrodynamics. Oxford: Oxford University Press, 222 Pp., \$49.95”. In: *Philosophy of Science* 74.4, pp. 555–558. DOI: 10.1086/524824.
- Norton, John D. (1993). “General Covariance and the Foundations of General Relativity: Eight Decades of Dispute”. In: *Reports of Progress in Physics* 56, pp. 791–861.
- Okon, Elias and Daniel Sudarsky (2016). “Less Decoherence and More Coherence in Quantum Gravity, Inflationary Cosmology and Elsewhere”. In: *Foundations of Physics* 46.7, pp. 852–879. DOI: 10.1007/s10701-016-0007-x.
- Penrose, R. (2016). *Fashion, Faith, and Fantasy in the New Physics of the Universe*. Princeton University Press. ISBN: 9781400880287. URL: <https://books.google.gr/books?id=T09kCwAAQBAJ>.
- Peskin, Michael E. and Daniel V. Schroeder (1995). *An introduction to quantum field theory*. Includes exercises. Boulder, CO: Westview. URL: <https://cds.cern.ch/record/257493>.
- Petrov, A.A. and A.E. Blechman (2015). *Effective Field Theories*. World Scientific Publishing Company Pte Limited. ISBN: 9789814434935. URL: <https://books.google.gr/books?id=tsDACwAAQBAJ>.
- Polo-Goméz, José, Luis J. Garay, and Eduardo Martin-Martinez (2021). “A detector-based measurement theory for quantum field theory”. In.
- Pooley, Oliver (2015). “Background Independence, Diffeomorphism Invariance, and the Meaning of Coordinates”. In: *Towards a Theory of Spacetime Theories*. Ed. by Dennis Lehmkuhl. Birkhäuser.
- Psillos, Stathis (1999). *Scientific Realism: How Science Tracks Truth*. Routledge.
- Reeh, H. and S. Schlieder (1961). “Bemerkungen zur unitäräquivalenz von lorentzinvarianten feldern”. In: *Il Nuovo Cimento* 22, pp. 1051–1068. DOI: 10.1007/BF02787889. URL: <https://doi.org/10.1007/BF02787889>.
- Rickles, Dean (2016). “Quantum gravity: A primer for philosophers”. In: *The Ashgate companion to contemporary philosophy of physics*. Routledge, pp. 268–388.

- Rivat, Sébastien (2019). “Renormalization Scrutinized”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 68, pp. 23–39. DOI: 10.1016/j.shpsb.2019.04.006.
- Rosaler, Joshua and Robert Harlander (2019). “Naturalness, Wilsonian Renormalization, and ?Fundamental Parameters? in Quantum Field Theory”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 66, pp. 118–134. DOI: 10.1016/j.shpsb.2018.12.003.
- Rovelli, Carlo (2007). *Quantum Gravity*. Cambridge Monographs on Mathematical Physics. Cambridge University Press. ISBN: 9781139456159. URL: <https://books.google.gr/books?id=16UgAwAAQBAJ>.
- Ruetsche, Laura (2011). *Interpreting Quantum Theories: The Art of the Possible*. Oxford University Press UK.
- (Feb. 2020a). “Perturbing Realism”. In: *Scientific Realism and the Quantum*. Oxford University Press, pp. 293–314. ISBN: 9780198814979. DOI: 10.1093/oso/9780198814979.003.0015.
- (Feb. 2020b). “Perturbing Realism”. In: pp. 293–314. ISBN: 9780198814979. DOI: 10.1093/oso/9780198814979.003.0015.
- (2018). “Renormalization Group Realism: The Ascent of Pessimism”. In: *Philosophy of Science* 85.5, pp. 1176–1189. DOI: 10.1086/699719.
- Rugh, Svend E. and Henrik Zinkernagel (2002). “The Quantum Vacuum and the Cosmological Constant Problem”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 33.4, pp. 663–705. DOI: 10.1016/s1355-2198(02)00033-3.
- Rugh, Svend E., Henrik Zinkernagel, and Tian Yu Cao (1999). “The Casimir Effect and the Interpretation of the Vacuum”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 30.1, pp. 111–139. DOI: 10.1016/s1355-2198(98)00034-3.
- Saatsi, Juha (Feb. 2020). “Truth vs. Progress Realism about Spin”. In: *Scientific Realism and the Quantum*, pp. 35–54. ISBN: 9780198814979. DOI: 10.1093/oso/9780198814979.003.0003.

- Sahuquillo, Miguel Ángel Carretero (2019). “The Charm Quark as a Naturalness Success”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 68, pp. 51–61. DOI: 10.1016/j.shpsb.2019.06.003.
- Salmon, Wesley C. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton University Press.
- Saunders, Simon (1993). “Decoherence, Relative States, and Evolutionary Adaptation”. In: *Foundations of Physics* 23.12, pp. 1553–1585. DOI: 10.1007/BF00732365.
- Schlosshauer Maximilian A., M.S.S.L.O. (2007). *Decoherence: And the Quantum-To-Classical Transition*. The Frontiers Collection. Springer. ISBN: 9783540357735. URL: <https://books.google.gr/books?id=1qrJUS5zNbEC>.
- Schwartz, Matthew D. (2014). *Quantum Field Theory and the Standard Model*. Quantum Field Theory and the Standard Model. Cambridge University Press. ISBN: 9781107034730. URL: <https://books.google.gr/books?id=HbdEAgAAQBAJ>.
- Schwinger, J. (1988). *Particles, Sources, And Fields (volume I)*. Advanced book classics. Basic Books. ISBN: 9780201094701. URL: <https://books.google.com/books?id=YcbvAAAAMAAJ>.
- Sen, Ashoke (2015). “Ultraviolet and Infrared Divergences in Superstring Theory”. In: *arXiv: High Energy Physics - Theory*.
- Signer, Adrian (May 2009). “ABC of SUSY”. In: *Journal of Physics G: Nuclear and Particle Physics* 36.7, p. 073002. ISSN: 1361-6471. DOI: 10.1088/0954-3899/36/7/073002. URL: <http://dx.doi.org/10.1088/0954-3899/36/7/073002>.
- Sklar, Lawrence (2000). *Theory and Truth: Philosophical Critique Within Foundational Science*. Oxford University Press.
- Smeenk, Christopher (2014). “Predictability Crisis in Early Universe Cosmology”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 46.1, pp. 122–133. DOI: 10.1016/j.shpsb.2013.11.003.
- Smeenk, Christopher and George Ellis (2017). “Philosophy of Cosmology”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2017. Metaphysics Research Lab, Stanford University.
- Srednicki, Mark (2007). *Quantum Field Theory*. Cambridge: Cambridge Univ. Press. URL: <https://cds.cern.ch/record/1019751>.

- Stachel, John (Jan. 2007). “The Story of Newstein or: Is Gravity Just Another Pretty Force?” In: pp. 1962–2000. DOI: 10.1007/978-1-4020-4000-9_46.
- Stanford, P. Kyle (2006). *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives: Science, History, and the Problem of Unconceived Alternatives*. Oxford University Press.
- Stein, Howard (1967). “Newtonian Space-Time”. In: *Texas Quarterly* 10.3, pp. 174–200.
- (2021). “Physics and Philosophy Meet: The Strange Case of Poincaré”. In: *Foundations of Physics* 51.3, pp. 1–24. DOI: 10.1007/s10701-021-00460-x.
- (1995). “Some reflections on the structure of our knowledge in physics”. In: *Logic, Methodology and Philosophy of Science IX*. Ed. by Dag Prawitz, Brian Skyrms, and Dag Westerstahl. Vol. 134. Studies in Logic and the Foundations of Mathematics. Elsevier, pp. 633–655. DOI: [https://doi.org/10.1016/S0049-237X\(06\)80067-4](https://doi.org/10.1016/S0049-237X(06)80067-4). URL: <https://www.sciencedirect.com/science/article/pii/S0049237X06800674>.
- (1992). “Was Carnap Entirely Wrong, After All?” In: *Synthese* 93.1-2, pp. 275–295. DOI: 10.1007/BF00869429.
- (2004). “Was Carnap entirely wrong, after all?” In: *Synthese* 93, pp. 275–295.
- (1989). “Yes, but... Some Skeptical Remarks on Realism and Anti-Realism”. In: *Dialectica* 43.1-2, pp. 47–65. DOI: 10.1111/j.1746-8361.1989.tb00930.x.
- Stoeltzner, Michael (Jan. 2004). “On Optimism and Opportunism in Applied Mathematics: Mark Wilson Meets John Von Neumann on Mathematical Ontology /at_lç”. In: *Erkenntnis* 60, pp. 121–145. DOI: 10.1023/B:ERKE.0000005144.79761.02.
- Suppes, Patrick Colonel (1957). *Introduction to Logic*. New York, NY, USA: Dover Publications.
- Susskind, Leonard (1979). “Dynamics of Spontaneous Symmetry Breaking in the Weinberg-Salam Theory”. In: *Phys. Rev. D* 20, pp. 2619–2625. DOI: 10.1103/PhysRevD.20.2619.
- (Mar. 2003). “The Anthropic Landscape of String Theory”. In: *arXiv: hep-th/0302219*. DOI: 10.1017/CBO9781107050990.018.
- (2008). *The Cosmic Landscape: String Theory and the Illusion of Intelligent Design*. Little, Brown. ISBN: 9780316055581. URL: <https://books.google.gr/books?id=RIW9E1sOyxUC>.

- Teller, Paul (1995). *An Interpretive Introduction to Quantum Field Theory*. Princeton University Press.
- Van Fraassen, Bas C. (1980). *The Scientific Image*. Oxford, England: Oxford University Press.
- Vickers, Peter (2013). *Understanding Inconsistent Science*. Oxford University Press.
- Wald, Robert M. (2018). “The Formulation of Quantum Field Theory in Curved Spacetime”. In: *Einstein Stud.* 14, pp. 439–449. DOI: 10.1007/978-1-4939-7708-6_15. arXiv: 0907.0416 [gr-qc].
- Wallace, David (2017). “Against Wavefunction Realism”.
- (2006). “In Defence of Naiveté: The Conceptual Status of Lagrangian Quantum Field Theory”. In: *Synthese* 151.1, pp. 33–80. DOI: 10.1007/s11229-004-6248-9.
- (2019). “Naturalness and Emergence”. In: *The Monist* 102.4, pp. 499–524. DOI: 10.1093/monist/onz022.
- (2020). “On the Plurality of Quantum Theories: Quantum Theory as a Framework, and its Implications for the Quantum Measurement Problem”. In: *Scientific Realism and the Quantum*. Ed. by S. French and J. Saatsi. Oxford University Press.
- (2010). “Taking Particle Physics Seriously: A Critique of the Algebraic Approach to Quantum Field Theory”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 42.2, pp. 116–125. DOI: 10.1016/j.shpsb.2010.12.001.
- (2012). *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford University Press.
- Weinberg, Steven (2007). “Living in the multiverse”. In: *Universe or Multiverse?* Ed. by Bernard Editor Carr. Cambridge University Press, pp. 29–42. DOI: 10.1017/CBO9781107050990.004.
- (1987). “Newtonianism, reductionism and the art of congressional testimony”. In: *Nature* 330, pp. 443–447. DOI: 10.1038/330433a0.
- (1995). *The Quantum Theory of Fields*. Vol. 1. Cambridge University Press. DOI: 10.1017/CBO9781139644167.
- (1996). *The Quantum Theory of Fields*. Vol. 2. Cambridge University Press. DOI: 10.1017/CBO9781139644174.

- Wells, James D. (2012). *Effective Theories in Physics: From Planetary Orbits to Elementary Particle Masses*. SpringerBriefs in Physics. Springer Berlin Heidelberg. ISBN: 9783642348921. URL: <https://books.google.gr/books?id=kKKAfrfCofcC>.
- (2019). “Naturalness, Extra-Empirical Theory Assessments, and the Implications of Skepticism”. In: *Foundations of Physics* 49.9, pp. 991–1010. DOI: 10.1007/s10701-018-0220-x.
- (2015). “The Utility of Naturalness, and How its Application to Quantum Electrodynamics Envisages the Standard Model and Higgs Boson”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 49, pp. 102–108. DOI: 10.1016/j.shpsb.2015.01.002.
- Wetterich, C. (1984). “Fine-tuning problem and the renormalization group”. In: *Physics Letters B* 140.3, pp. 215–222. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(84\)90923-7](https://doi.org/10.1016/0370-2693(84)90923-7). URL: <https://www.sciencedirect.com/science/article/pii/0370269384909237>.
- Williams, Porter (2015). “Naturalness, the Autonomy of Scales, and the 125GeV Higgs”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 51, pp. 82–96. DOI: 10.1016/j.shpsb.2015.05.003.
- (2019a). “Scientific Realism Made Effective”. In: *British Journal for the Philosophy of Science* 70.1, pp. 209–237. DOI: 10.1093/bjps/axx043.
- (2019b). “Two Notions of Naturalness”. In: *Foundations of Physics* 49.9, pp. 1022–1050. DOI: 10.1007/s10701-018-0229-1.
- Wilson, Kenneth G. (1975). “The Renormalization Group: Critical Phenomena and the Kondo Problem”. In: *Rev. Mod. Phys.* 47, p. 773. DOI: 10.1103/RevModPhys.47.773.
- Wilson, Kenneth G. and J. Kogut (1974). “The renormalization group and the expansion”. In: *Physics Reports* 12.2, pp. 75–199. ISSN: 0370-1573. DOI: [https://doi.org/10.1016/0370-1573\(74\)90023-4](https://doi.org/10.1016/0370-1573(74)90023-4). URL: <https://www.sciencedirect.com/science/article/pii/0370157374900234>.
- Wilson, Mark (2017). *Physics Avoidance: And Other Essays in Conceptual Strategy*. Oxford University Press.
- (2006). *Wandering Significance: An Essay on Conceptual Behavior*. Oxford: Clarendon Press.

- Wimsatt, William C. (2007). *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Harvard University Press.
- Winther, Rasmus Grønfeldt (2021). “The Structure of Scientific Theories”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2021. Metaphysics Research Lab, Stanford University.
- Witten, Edward (2015). “What every physicist should know about string theory”. In: *Phys. Today* 68.11, pp. 38–43. DOI: 10.1063/PT.3.2980.
- Woit, Peter (2006). *Not Even Wrong: The Failure of String Theory and the Search for Unity in Physical Law*. Basic Books. ISBN: 9780465092758. URL: <https://books.google.gr/books?id=pcJA3i0xKAUC>.
- Zelko, Martin N. (forthcoming). “Explaining Thermodynamics: Doctoral Dissertation”.