Numerical Investigations of the Dirac Equation and Bound State Quantum Electrodynamics in Atoms

A thesis presented in partial fulfilment of the
requirements for the degree of

Doctor of Philosophy
in
Physics

at
Massey University, Albany
New Zealand

Morten Piibeleht

2022

Abstract

This thesis addresses, from a computational perspective, several open questions in relativistic atomic structure theory, which is the theoretical description of atoms based on the Dirac equation and quantum electrodynamics (QED).

The first part of this thesis investigates several fundamental problems of the Dirac equation with the help of a novel numerical solver based on the one-dimensional finite element (FEM) basis set. Significant effort is made to validate and benchmark the solver, which is reliably able to converge to accurate results at numerical floating-point precision, including when nuclear potentials derived from nuclear models with finite spacial extent (as opposed to a point nucleus) are used. The solver is then applied to the Dirac equation in the challenging high nuclear charge regime where the Dirac equation exhibits several mathematical difficulties. In particular, the problem of the $1s$ bound state diving into the sea of negative energy continuum states is studied and the diving resonance state is numerically traced and analysed. As a type of workaround, a modified version of the Dirac equation where the negative energy plane-wave states are projected out of the Hilbert space is also solved and studied in the high nuclear charge regime.

The second part of the thesis involves expanding the QED self-energy treatment in the atomic structure software GRASP. The configuration interaction (CI) portion of the code is significantly refactored to allow for the implementation of new additional effective operators that provide a more modern multi-electron treatment of QED self-energy effects. The implementation is tested by evaluating the QED and other post-Dirac-Coulomb corrections for the ground states of the beryllium-like isoelectronic sequence, which was also discovered to exhibit an interesting ground state configuration transition at high nuclear charge.

Acknowledgements

It's been a journey, aye? While I personally find the acknowledgements section of any thesis a bit awkward, I nevertheless do want to mention the people who, in one way or another, have supported or contributed to this thesis. So here goes.

First, Peter. Thank you for supervising this work, all your support, and having my back. It has been a privilege to work here, bringing plenty of great experiences, and allowing me to continue learning about physics in a field that genuinely captures my imagination. And also Patrick, thank you for heaps of groovy chats on these topics!

On the research side, I wish to mention and thank the people in the CompAS collaboration — Jon, Per, Jörgen, Henrik, Charlotte, Ian, Gediminas, Michel, Tomas, Kai, Ran, Jacek, to name a few — for a truly friendly collaborative environment, hosting my visits, memorable collaboration workshops, and for the many hours of insightful conversations and emails about GRASP. Perhaps another visit to the Malmö Symphony is also in order at some point. I should also mention that it was also a great fun to run into some old friends from Lund (Madeleine, Asimina, Brian) and to visit the old stomping grounds in Skåne.

I also want to thank Susi, for assisting in getting the FEM project off the ground, and all the work and help with the HelFEM code base. In the same vein, thank you Paul for hosting me in Paris, and for the discussions on QED and the atomic structure methods.

Back in New Zealand: Vesna, for the printing accounts, other paperwork, and helping me navigate the Massey bureaucracy. Mike, for being there with any computer cluster or other IT questions. Margaret and Sam, for the distraction involving teaching undergrads and shooting bottle rockets filled with liquid nitrogen. And, finally, all the great friends from the CTCP who made this journey fun, in no particular order: Odile, Peter, Antony, Ray, Sarthak, Lukas, Lukas and Lukaš, Edison, Antonio, Ulrich, Jan, Jan, Steffi, Lotta, Ronja, Sophie, Jayson, Matija, Joachim, Elke. A big thank you to all of you for your company.

Ja lõpuks kaugel kodumaal — aitäh Pille, Memm ja kõik head sõbrad, et olete ikka seal, mis siis, et mul ainult harva võimalik külas käia oli, eriti doktori viimasel kahel aastal.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

An atom seems like one of the simplest systems in physics: a positively charged atomic nucleus in the middle with a flock of electrons surrounding it. Physics students solve it with quantum theory in an undergraduate course. Job done. Right? No?

The problem, if it can be called that, is precision. The advances in the atomic spectroscopy of neutral atoms and charged ions allow for extremely accurate measurements of transitions energies, intensities, and other associated properties. This precision means that atomic systems are useful tools for other areas of physics, from helping to determine the atmospheric composition of stars and exoplanets and understanding the behaviour of plasma in astrophysics or fusion reactors [1] to testing fundamental physics theories [2, 3, 4].

It is also important for theory to keep up with this progress. And, by and large, it has. Atomic structure theory has come a long way since its early days in the form of highly empirical formulae for hydrogen spectral lines proposed by Balmer and Rydberg in the late 1880s. Via early developments in quantum theory in the form of Bohr's atomic model and then non-relativistic Schrödinger equation, to the relativistic generalization by Dirac, the theoretical understanding of atoms, and elementary particle more generally, has arrived at quantum field theories (QFT). It was the desire to properly describe atoms that spurred Feynman, Schwinger, Dyson, and many others to develop quantum electrodynamics (QED), the first fully relativistic and quantized theory of electrons and their electromagnetic interactions. As new particles were discovered or proposed — quarks, neutrinos, the Higgs boson — the same methods were then applied to those, leading to the Standard Model of Particle Physics, a quantum field theory of all elementary particles.

As a theory, the Standard Model is very stringently tested, and QED in particular allows for extremely precise theoretical predictions — already in 1985, Feynman himself called QED the "jewel of physics"[1]. As an example, the QED contributions to the electron magnetic moment $g - 2$, and hence the fine structure constant, has been determined to ten significant figures by summing over twelve thousand tenth order Feynman diagrams [6], with experiments achieving similar precision [7]. Much of the work nowadays is to experimentally put bounds of various extensions or generalizations of the Standard Model (e.g. string theory, supersymmetry, grand unified theories).

Nevertheless, all that progress does not mean that there are no challenges remaining. Even the basic mathematical formalism of QFT has plenty of significant open questions, such as the divergences that make it difficult to arrive at finite results in quantum field theory. While for most practical purposes these divergences can be resolved via renormalization group theory and by

[1] *"So far we have found nothing wrong with QED. It is, therefore, the jewel of physics — our proudest possession."* in "QED: The Strange Theory of Light and Matter" [5].

treating existing quantum field theoretic models as effective low-energy theories [8, Chapter III.1], that does not actually remove the divergences from the underlying mathematical model. How to fundamentally get rid of the divergences has not been solved in the general case [9, Preface & Chapter 8]. In a similar vein, Haag's theorem [10] and its generalization show that the interaction picture, which forms the very basis of the perturbative treatment of quantum field theories, is inconsistent, and in fact the current axiomatic approaches to QFT are inconsistent with quantum field theories involving interacting fields [9, Chapter 8].

In relativistic quantum theory, the non-field theoretic single-particle Dirac equation has open questions, especially in the high nuclear charge regime. It is well-known that, for a point nucleus, the solutions of the Dirac equation break down at $Z = \alpha^{-1} \approx 137$, evident from the Sommerfeld formula for the bound state energies

$$E_{n\kappa}(Z) = mc^2 \left(1 + \frac{(Z\alpha)^2}{\left[n - |\kappa| + \sqrt{\kappa^2 - (Z\alpha)^2} \right]^2} \right)^{-1/2}$$

predicting imaginary energies if $Z > \alpha^{-1}$, and is related to the Dirac operator losing its self-adjointness in that regime [11]. Back in the day, Feynman stated that "137 is one of the greatest damn mysteries of physics: a magic number that comes to us with no understanding" and believed that nuclear charges beyond this number do not make sense physically [12]. While we know that for a more physical nuclear model with a finite extent where the charge is smeared over a small but finitely sized volume this is no longer an issue, the fact that the Dirac equation exhibits negative continuum states nevertheless causes problems. At a high enough nuclear charge, bound states will dive into the negative continuum, leaving the physical interpretation unclear.

These issues with the single-particle Dirac equation also have practical and interpretational implications for QED and many-electron theory. In the Furry picture of bound-state quantum electrodynamics, the reference states are exactly the solutions to the single-particle Dirac equation with a nuclear potential. Therefore, if the reference solutions break down, so does the whole procedure. Similarly, numerical solvers of the single-particle equations of single- or multiconfigurational Dirac-Hartree-Fock procedures generally break down when the energies of the states drop below $-mc^2$. In both cases, in order to perform calculations in those regimes, it is important to understand how to properly treat those states. While solutions, such as tightening up the boundary conditions [13, 14] or treating the system in a *rigged Hilbert space* formalism, leading to Gamow states [15], have been proposed, the correct treatment of such diving states is still up for debate.

A problem of a more practical nature has to do with the computational complexity of large quantum systems. The quantum entanglement, or correlation, of particles means that in general there is a combinatorial explosion with increasing number of particles as to how much information a system can contain, and therefore how much memory and computational time is required for its study. This is further exacerbated if full interacting QED is taken into account where particles can also be created and annihilated. There are methods, such as configuration interaction, many-body perturbation theory or even density functional theory, that allow us to get very good results for even relatively large atomic systems, such as the super-heavy elements, but these methods do not allow for a full QED treatment of the systems. On the other hand, a full QED treatment of a many-electron system is extremely expensive computationally, and even performing calculations with four electrons (a Be-like system) is considered to be a significant achievement [16, 17]. However, it is possible to extend the aforementioned methods to approximately account for QED with effective potentials [18, 19, 20], where the error from treating QED approximately is

expected to be less than the error arising from the inability to fully account for the many-electron correlation.

Overview of the thesis

This thesis is concerned with the quantum electrodynamics (QED) sector of the Standard Model. The goal is to advance the methods for highly accurate theoretical studies of atoms, especially in the challenging high nuclear charge region, where even the relatively simple Dirac equation runs into trouble. The error made by disregarding QED when studying electronic systems is relatively small, but it is still significant and experimentally resolvable. The error from disregarding QED also increases with increasing nuclear charge [21], and so it is important to consider QED when studying heavy and super-heavy elements.

With those challenges in mind, the work in this thesis makes two main contributions. First, is the numerical study of the Dirac equation in the high nuclear charge regime using the finite element method, with the goal of systematically benchmarking the method and exploring this challenging regime. The second major part is the implementation of additional QED operators in the configuration interaction framework of the GRASP software [22], which also involves significant changes to the CI program, in order to properly test the implementation and make future maintenance easier. This lays the groundwork for accurate computations of multi-electron (including open-shell) systems that can accurately include QED contribution with the GRASP software.

The core chapters of the thesis are laid out as follows:

- Chapter 2 gives a concise overview of the theoretical basis necessary for understanding this thesis. The aim is to introduce all the terminology from the ground up. The starting point is the Standard Model of particle physics which is then simplified and reduced to the models and equations used in practical calculations.

- Chapter 3 focuses on demonstrating how, based on a finite element basis set, it is possible to build a simple numerical solver that can be used to probe the relativistic Dirac equation in the high-Z regime, where its proper treatment becomes unclear. The solver can, to machine precision, solve the Dirac equation and makes it easy to interpret the solutions in various ways. It can also be extended very easily to solve related cases, which is demonstrated by the study of the energy-projected Dirac equation.

- Chapter 4 focuses on the treatment of QED effects in many-electron methods, specifically in the multiconfigurational Dirac-Hartree-Fock and configuration interaction (CI) method implemented in the GRASP software. In particular, the focus is on how to extend GRASP with additional QED self-energy operators, and then the implementation is tested on a Be-like isoelectronic sequence.

- Finally, Chapter 5 makes a few concluding remarks, discussing the work in a wider context and laying out a few avenues for future work. In addition, the thesis contains three appendices, with Appendix A and Appendix B functioning as a reference for a few basic mathematical results relevant for this work, and Appendix C containing a few extra figures for Chapter 3.

Units and notation

Most expressions, unless stated otherwise, are given in the "particle physics" natural units, where the reduced Planck constant $\hbar$, the electron mass m_e, the speed of light c, and the vacuum permittivity ε_0 are set to unity, i.e.

$$\hbar = m_\mathrm{e} = c = \varepsilon_0 = 1$$

That said, sometimes it is still convenient to ignore that those quantities are formally unity and still say e.g. that "energy is mc^2", rather than "energy is 1", for the sake of clarity.

However, another set of natural units that is occasionally used (mostly in Chapter 4, but also briefly discussed in Section 2.6) are the Hartree atomic units, where in addition to m_e and $\hbar$ it is the elementary charge e and Coulomb's constant $k_\mathrm{e} = \frac{1}{4\pi\varepsilon_0}$ that are set to unity

$$\hbar = m_\mathrm{e} = e = k_\mathrm{e} = 1$$

In this system, distances are measured in Bohr radii a_0, and energies are measured in Hartrees E_h. In this system, the speed of light is $c = 1/\alpha \approx 137$, where α is the dimensionless fine structure constant [23]

$$\alpha = \frac{1}{4\pi\varepsilon_0}\frac{e^2}{\hbar c} = 0.0072973525693(11)$$

Consequently, the rest mass of the electron is $m_\mathrm{e}c^2 \approx 18778\ E_\mathrm{h}$.

Finally, Appendix A functions as a repository for various conventions and mathematical notation that is used throughout this thesis.

Chapter 2

Elements of atomic theory

While the core parts of this thesis do not directly deal with quantum field theories (QFTs), it is important to recognize that all the underlying theory can, in principle, be derived from a single mathematical expression: the Lagrangian of the Standard Model of Particle Physics. This is worth pointing out, as the two main concepts that are focussed on in this thesis — the single-particle Dirac equation and relativistic many-body theory — are often treated in a self-contained way in the literature, without directly referring to the underlying quantum field theory.

This section gives a concise overview of the necessary theoretical background and notation. Rather than following the history chronologically, or being too concerned about pedagogical clarity and simplicity[1], it is organized in a top-down fashion, starting from the aforementioned Standard Model, and then, step by step, zooming into the corner of the theory that we are actually interested in. This allows the complexity to be systematically reduced, with a reasonably rigorous justification for each step. The chapter does not contain any original research *per se*, other than perhaps some interpretational points, and lays the groundwork for the subsequent chapters instead.

2.1 The Standard Model

As of now, the Standard Model of Particle Physics is currently the most precise and experimentally verified description of interacting microscopic elementary particles. In principle, every calculation performed in atomic physics is derived from that theory, even if this is not immediately apparent when dealing with standard relativistic quantum mechanics. It is all a matter of making reasonable approximations of the complete Standard Model Lagrangian to make calculations more tractable.

Particles

As a quantum field theory, the Standard Model describes particles as quantum fields, which can mathematically be defined as operator-valued functions over space and time. All the fields and their interactions are combined into a single mathematical object: the *Lagrangian density* $\mathcal{L}$[2]. This mathematical expression, analogous to the Lagrangian of a classical system, contains all the physics of the theory. One thing to note here is that this does not mean that there is a

[1] For the latter, there are many excellent textbooks and lecture series available, some of which will be mentioned in future footnotes.

[2] Giving the exact mathematic expression for $\mathcal{L}$ would not serve much purpose here as in its full glory, it is quite lengthy, would require a deeper discussion, and can be found anyway in standard textbooks [24, 8, 25].

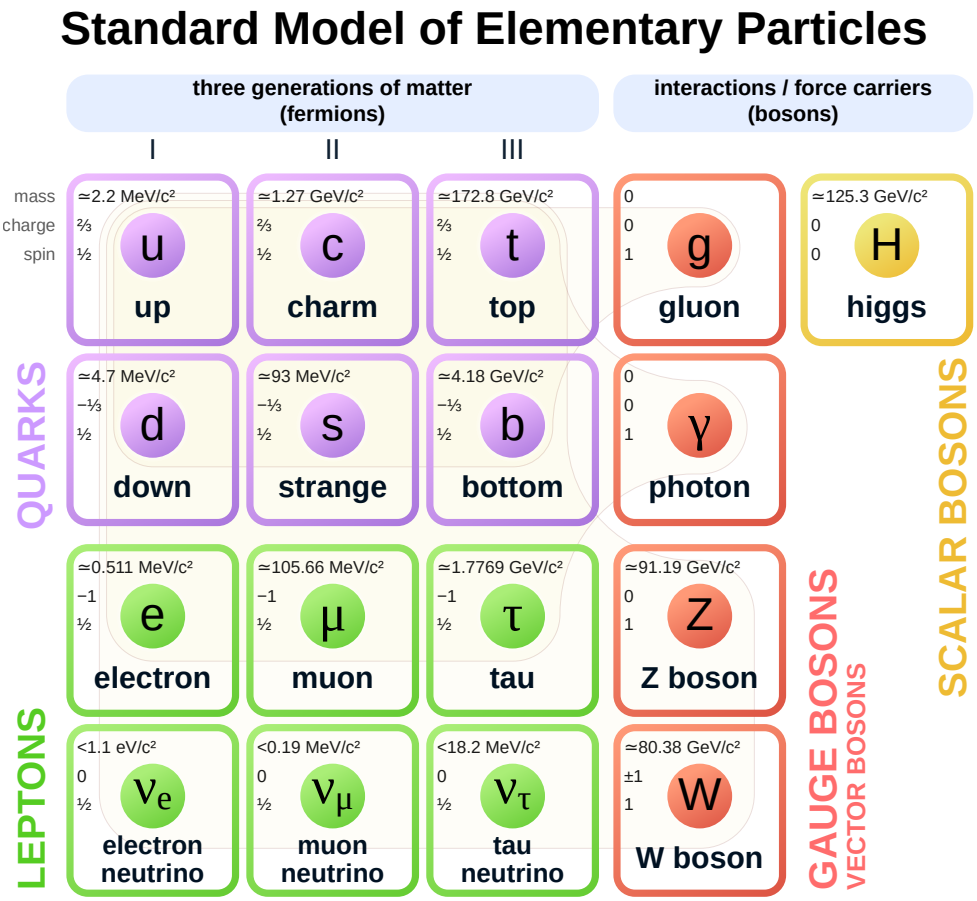

Figure 2.1: The "periodic table" of both the matter and interaction particles in the Standard Model, illustrating how they can be grouped in various ways. For each particle, their three most fundamental properties are listed — mass, charge, and spin. *Attribution: Cush / Wikimedia Commons / Public Domain, with masses updated to the latest PDG values [26].*

quantum field for each particle in the classical sense. Rather each type of particle (e.g. electron or photon) is described by a field, and individual particles in the classical sense would essentially be different excited states of the quantum field with different particle numbers. The Standard Model is also inherently relativistic, meaning that it takes into account the symmetries of special relativity, which is manifested in the Lorentz invariance of the Lagrangian density.

The particles, or fields, included in the Standard Model can be organized into a type of "periodic table", as visualized in Figure 2.1. According to their inherent behaviour under rotational symmetry transformations, quantified by the *spin* quantum number, they split into two categories: (1) fermionic matter particles with half-integer spin of 1/2, obeying the Pauli principle, and (2) the bosonic interaction fields with integer spin.[3] The spin 1/2 of fermions, which can

[3]When studying the statistical mechanics of many microscopic particles, then this difference in the fundamental quantum properties of the particles leads to different statistical distribution, referred to as the *Fermi-Dirac*

be interpreted as taking two full rotations of space to return a fermion back into its original state, has the consequence of not allowing two particles (in the classical sense) to occupy the same state. The interaction bosons mostly are spin-1 *vector particles*, essentially behaving like a standard three-vector under spacial rotations, except for the Higgs boson, which is a spin-0 *scalar* particle.

Fermionic matter particles can further be split into three "generations" — essentially the same particle with respect to most physics attributes, but with increasing masses. Only the first generation is what we deal with in most day-to-day, even in atomic physics, with the II and III generations being unstable and requiring high-energy processes to be created. On the other hand, the fermions split into two groups: quarks (fractionally charged particles, first generation of which make up the nuclei of atoms) and leptons (containing the charged electron and its higher-generation counterparts, and the almost interactionless, chargeless and almost massless neutrinos).

Interactions

As was mentioned, the Standard Model particles are either fermionic matter particles or vector bosons that mediate interactions. The classical example, specifically relevant to atomic physics, is the electromagnetic interaction, mediated by the photon, influencing charged matter particles like the electron and the muon.

The mathematical description of the interactions in quantum field theories brings in another key concept — the *gauge symmetry*. That is, each interaction can be associated with a unitary group, and the full theory (i.e. the Lagrangian density) is expected to be invariant under that group action. The conserved density associated with the gauge symmetry transformation is the charge of the interaction, such as the electric charge or the weak charge. In terse mathematical notation, the Standard Model is therefore often summarized as a

$$SU(3) \times SU(2) \times U(1)$$

gauge theory. The different terms correspond to the different interactions present: $U(1)$ for electromagnetic, $SU(2)$ for the weak interaction, and $SU(3)$ for the strong interaction.

Limitations of the Standard Model

While the Standard Model is extremely good at describing elementary particles, we should not shy away from mentioning its shortcomings. In addition to philosophical criticisms of the model being fine-tuned and containing too many parameters, for example, it is also by no means a *theory of everything*. [27, 28] The following lists a few examples where the Standard Model either does not even attempt to explain some observed physics, or gets it wrong.

Gravity. A glaring omission in the Standard Model is the lack of gravity. Heuristically, this is not a problem, as the weakness of the gravitational interaction relative to the strong, weak and electromagnetic interactions means that it can pretty much be ignored on the level of elementary particles[4] However, there is no well-accepted approach to even describe gravity together with the quantum description of elementary particles, as all attempts to consistently include the gravitational force in the Standard Model are riddled with theoretical and mathematical problems [29, 30]. In fact, in 2011 it was noted by Sushkov et al. that "It is remarkable that two of

statistics and *Bose-Einstein statistics* for fermions and bosons, respectively.

[4]The ratio of the electromagnetic vs. the gravitational force between the two charged particles electron and proton is $e^2/Gm_em_p = 2.26881776 \times 10^{39}$.

the greatest successes of 20th-century physics, general relativity and the standard model, appear to be fundamentally incompatible" [31].

Neutrino masses. A crucial property of a neutrino is that only the left-handed component interacts with other particles via the weak interaction, with the right-handed component remaining inactive. In order to incorporate this into the Standard Model in a Lorentz-covariant way, while still requiring the neutrinos to be Dirac fermions, it is possible to assume the neutrinos to be massless. However, the observation of neutrino oscillations — that neutrinos, as they propagate through space and time, can change their flavour, e.g. transition from electron neutrino to a muon neutrino — is at odds with this. To explain these oscillations, it is necessary to include neutrino masses and mixing angles in the model, such that their mass eigenstates become slightly rotated relative to the flavour eigenstates. As such, the exact theoretical nature and even the absolute hierarchy of masses of the neutrinos is still unknown. [32]

Dark matter. Astrophysical and cosmological observations of galaxies, galaxy clusters and the structure of the universe at large have long implied the existence of unexplained mass in the universe, required to explain how various structures have formed, and the motion of stars and galaxies. Explanations supposing that this may be due to hard-to-detect matter, such as brown dwarfs, neutron stars, black holes, or unassociated planets, have to large degree been ruled out by astrophysical surveys. While there are models attempting to explain dark matter via modifications to the gravitational interactions, arguably the most widely held assumption is that this dark matter is due to some unknown, very weakly interacting particles that, often referred to as WIMPs (Weakly Interacting Massive Particles). [33, 34]

Dark energy. The cosmological constant in Einstein's field equations in the general theory of relativity can be interpreted as the energy that drives the accelerating expansion of the universe. It has been suggested that dark energy could be explained as the energy density of vacuum fluctuations of the quantum fields, but back of the envelope estimates of such fluctuations are hundreds of orders of magnitude larger than the observed dark energy density. [35, 36] As such, the Standard Model currently does not provide a solid explanation for dark energy.

Anomalous magnetic moment of the muon. A long-standing discrepancy between experiments and theoretical calculations based on the Standard Model is in the value of the anomalous magnetic moment $a_\mu = \frac{g_\mu - 2}{2}$ for the muon. When considering only the Dirac equation, the expected value for the magnetic moment of a lepton is exactly $g = 2$. But as has been seen in experiments, and can be shown theoretically by considering the full Standard Model, including contributions from the strong and weak interactions, the value actually deviates from 2, leading to a non-zero value for a_μ [37]. However, as of 2021, with a significant 4.2σ deviation, state-of-the-art theoretical calculations are still not able to reproduce the experimental value latest experimental results [38]. While it can not be ruled out with absolute certainty that future improvements to either theoretical calculations or to experiments will not resolve this discrepancy, it may also be an indicator of undiscovered physics not included in the Standard Model.

2.2 The fields of Quantum Electrodynamics

Not all parts of the Standard Model are important when the intent is to describe the behaviour of electrons in atoms and molecules to a certain accuracy. Rather, it is sufficient to focus on what is commonly referred to as the quantum electrodynamics or QED sector of the Standard

Model. Compared to the full theory, only electrons[5], photons, and their interactions are retained. Atomic nuclei are usually reduced to an effective charge distribution, as the fact that they are actually dynamical systems made up of strongly interacting quarks can be, to a very good approximation, ignored. Additional corrections, arising from the interactions with the other particles of the Standard Model such as the weak bosons, can always be included back in the model, possibly also in some approximate way [39].

As a historic note, the QED sector of the Standard Model is really the original quantum field theory, developed in the 1930s and 1940s. [40, 41, 42] It arose out from the need to extend the relativistic description of electrons (i.e. the Dirac equation) by also including the quantized nature of the electromagnetic interaction. One of the original experimental motivations was the Lamb shift [43], where it became clear there is an unexplained energy splitting in the spectra of the hydrogen which is not explained by the Dirac equation, even if vacuum polarization is taken into account [44, 45].

Particles, fields and Lagrangians

The Lagrangian density of QED contains only two fields: the 4-component spin-1/2 fermionic charged lepton field ψ, simultaneously describing electrons and positrons, and the electromagnetic field A^μ describing photons. The other fields of the Standard Model are removed and what is left is the expression

$$\mathcal{L} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + i\bar{\psi}(\partial_\mu\gamma^\mu - m)\psi - e\bar{\psi}\gamma^\mu\psi A_\mu + j^\mu_{\text{ext}}A_\mu$$

where $F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu$ is the electromagnetic field tensor and $\bar{\psi} = \psi^\dagger\gamma^0$ the Dirac adjoint.[6] The different terms in the Lagrangian density can be given physical interpretations as well:

- $-\frac{1}{4}F^{\mu\nu}F_{\mu\nu}$ describes the dynamics of free photons or, equivalently, the quantized electromagnetic field;
- $i\bar{\psi}(\partial_\mu\gamma^\mu - m)\psi$ describes free electrons and positrons (or, more generally, a charged lepton with mass m);
- $-e\bar{\psi}\gamma^\mu\psi A_\mu$ is an *interaction term* between the two fields, and describes how photons and electrons affect each other, with the elementary charge e determining the strength of the interaction;
- Lastly, $j^\mu_{\text{ext}}A_\mu$ is the interaction term of the quantized electromagnetic field with an external charge distribution $j^\mu_{\text{ext}}(x)$. It is worth pointing out that we can also calculate the four-current for the quantized electrons, with $j^\mu = e\bar{\psi}\gamma^\mu\psi$, which makes it clear that the last two terms have an identical mathematical structure.

In principle, the external charge-current density j^μ_{ext}[7], is not part of the Standard Model Lagrangian density. However, adding this into the QED one allows us to incorporate nuclear charges and other sources of external EM fields into our models. As such, j^μ_{ext} depends on the particular system under study, whereas the rest of the Lagrangian density is universal.

[5] In some cases, the other charged leptons, usually muons, may also be included. Mathematically, their description is the same as for the electron, requiring a spin-1/2 Dirac field, just with a higher mass.

[6] For the definition of the four-by-four Dirac gamma matrices, see Appendix A.

[7] The time- and space-like components of a charge-density four-current are the charge density ρ and the vector 3-current $\boldsymbol{j}$, respectively, i.e. $j^\mu = (\rho, \boldsymbol{j})$. All of them are fields over space and time.

If it is necessary to add other leptons, such as the muon, to the model, it is simply a matter of defining an additional Dirac field, and adding the corresponding free-particle term with the appropriate mass together with an interaction term for that new Dirac field into the Lagrangian density. It should be noted that calculations with muons are very relevant, as due to their heavy mass they stay much closer to the nucleus, are therefore very useful for probing the size and shape of nuclei [46].

Quantum field theory and Fock spaces

Simply writing down the Lagrangian density of a theory is not enough to get quantitative results. Now, especially for time-dependent processes, the *path integral* formulation of quantum field theories can be very elegant and convenient [47, 42, 8]. However, as we are interested in bound states in atomic systems, it will be easier to follow the more common and traditional so-called *second quantization* formulation [24], which leads to the familiar Hamiltonian operator based quantum mechanics.

The proper mathematical formulation of second quantization, which gives a more precise definition for the field operators, relies on the concept of a Fock space[8]. [48] The construction of a Fock space starts by considering a single-particle Hilbert space $\mathcal{H}$ of a field (or, equivalently, of a particle species), which we further assume is spanned by some complete set of basis functions $\{\psi_i(\boldsymbol{x})\} \subset \mathcal{H}$. This Hilbert space is then generalized to a many-particle space.

For electrons-positrons, which are fermions and therefore anti-symmetric under particle interchange, an N-particle space is constructed via an N-fold anti-symmetric tensor product of the single-particle Hilbert space $\wedge^N(\mathcal{H})$. The N-fold anti-symmetric tensor products of the basis states of the Hilbert space $\psi_i(\boldsymbol{x})$ form a basis for the N-particle space $\wedge^N(\mathcal{H})$. Those spaces, with a fixed number of particles, are then added together via an infinite direct sum

$$\mathcal{F}_{\mathrm{f}}(\mathcal{H}) = \bigoplus_{N=0}^{\infty} \wedge^N(\mathcal{H})$$

to form the full fermion Fock space over the underlying Hilbert space. Physically, the elements of this space represent wavefunctions with any number of particles and, in fact, states that are superpositions of states with different numbers of particles.

In the case of bosons, such as when constructing the boson Fock space of photons, the logic is similar, but one must instead employ the N-fold *symmetric* tensor product $\otimes_{\mathrm{s}}^N$ when constructing the N-particle subspaces. The direct sum

$$\mathcal{F}_{\mathrm{b}}(\mathcal{H}) = \bigoplus_{N=0}^{\infty} \otimes_{\mathrm{s}}^N(\mathcal{H})$$

again leads to the full boson Fock space over the underlying Hilbert space.[9]

In the so-called second quantization notation, more familiar to readers of quantum mechanics and quantum field theory literature, states in the N-particle space can be expressed with the help of creation and annihilation operators. Specifically, we assume that each of the basis states $\psi_i(\boldsymbol{x})$ of $\mathcal{H}$, which we also assume are orthogonal to each other, has a corresponding creation operator

[8] An excellent review of the mathematical underpinnings of Fock spaces can be found in Arai's 2018 book on the subject. [9]

[9] The symmetric and anti-symmetric tensor products can be defined as a subset of the standard, full N-fold tensor product of the Hilbert space, i.e. $\otimes_{\mathrm{s}}^N(\mathcal{H}P) \subset \otimes^N(\mathcal{H})$ or $\wedge^N(\mathcal{H}P) \subset \otimes^N(\mathcal{H})$, of states with the desired symmetry or anti-symmetry property. The fermion and boson spaces are disjoint subsets of the full tensor product space. Section 2.9.4 in Arai's book offers more detailed definitions [9].

$c_i^\dagger$ and annihilation operator c_i associated with it, with the former being the adjoint of the latter. Respectively, they have the effect of adding or removing a particle in that particular state to or from the many-particle state. To ensure the appropriate symmetry properties arising from indistinguishability, the creation and annihilation operators obey the commutation relations

$$[c_i, c_j^\dagger] = \delta_{ij}, \quad [c_i, c_j] = 0$$

in the case of bosons, and anti-commutation relations

$$\{c_i, c_j^\dagger\} = \delta_{ij}, \quad \{c_i, c_j\} = 0$$

in the case of fermions.

The creation and annihilation operators are very much analogous to the ladder operators in the single-particle quantum harmonic oscillator theory. For both fermions and bosons, trying to remove a particle that does not exists in a state leads to the many-particle state being fully annihilated (i.e. $c_i |\Psi\rangle = 0$). However, for fermions, the anti-commutation relations further have the effect that, when adding a particle to a state that already contains that single particle state, the many-particle states again is fully annihilated (i.e. $c^\dagger |\Psi\rangle = 0$). This is not the case for bosons, and in other words this inherently takes into account the Pauli principle, where multiple particles can not be in the same state at once.

Physically interesting states are then constructed by adding particles to the *vacuum state* $|0\rangle$ — the only element of the 0-particle space $\wedge^0(\mathcal{H}) = \{|0\rangle\}$.[10] The element of the Fock space that corresponds to a single-particle basis function $\psi_i(\boldsymbol{x})$ can be written as

$$|i\rangle = c_i^\dagger |0\rangle$$

A basis for a full N particle space can be generated by considering the products on N creation operators applied to the vacuum state

$$|i_1 i_2 \dots i_N\rangle = c^\dagger(i_1) \cdot c^\dagger(i_2) \cdots c^\dagger(i_N) |0\rangle$$

It is necessary to take into account that the indices i_n can not repeat and are chosen in a way that the same combination does not appear twice. The latter can be achieved by requiring them to be ordered such that $i_1 < i_2 < \cdots < i_N$.

Overall, given the Hilbert space $\mathcal{H}_e$ of the electron-positron field, and $\mathcal{H}_\gamma$ of the photon field, the full QED Fock space is simply constructed as a tensor product of those space

$$\mathcal{F}_\mathrm{b}(\mathcal{H}_\gamma) \otimes \mathcal{F}_\mathrm{f}(\mathcal{H}_e)$$

In case there are additional fermions in the theory, such as muons, bringing their own Hilbert space $\mathcal{H}_\mu$, the product expands to

$$\mathcal{F}_\mathrm{b}(\mathcal{H}_\gamma) \otimes \mathcal{F}_\mathrm{f}(\mathcal{H}_e) \otimes \mathcal{F}_\mathrm{f}(\mathcal{H}_\mu)$$

[10] Formally, as the vacuum state by definition contains no particles, it can also be defined as a state that gets annihilated whenever a particle is removed, i.e.

$$\forall k : \hat{c}_k |0\rangle = 0$$

Field decomposition

In order to connect the Lagrangian of a theory to the Fock space formalism, it is necessary to give a precise definition for the quantum field operator $\Psi(\boldsymbol{x})$, which is an operator acting on the Fock space. Given the basis states $\psi_i(\boldsymbol{x})$ and the corresponding single-particle creation and annihilation operators $c_i^\dagger$ and c_i, the field operator can be written as

$$\Psi(\boldsymbol{x}) = \sum_i \psi_i(\boldsymbol{x}) c_i$$

where the sum runs over all the states. This actually represents a whole set of operators, a different one for each point in space, and the argument $\boldsymbol{x}$ can be thought of as a label. Like the annihilation operators it is made up from, such a field operator also has the effect of annihilating a particle from a state.

Giving the field operator a precise definition like this enables quantitative calculations to be performed within the Fock space formalism. At the same time, even though it is an operator, the field operator can be used very much like a multi-variable function, allowing the Lagrangian densities and Hamiltonians to be writted down in a more physical form with derivatives etc.

The only remaining question is the exact choice of basis states, and most often the field decomposition is written down in terms of the free particle solutions of the field. For example, for the electron field, the starting point would be the non-interacting, unbound free-electron part of the QED Lagrangian density. From that, by applying the variational principle, it is possible to derive the equations of motions for a free particle, which for the electron would be the free particle Dirac equation. The solutions to those equations, which generally take the form of plane waves and can usually be solved for via a simple application of the Fourier transform to the partial differential equation, are then used in expanding the field operator. This procedure for various fields is described in standard textbooks [24, 25, 49] and the following section just reviews the Dirac field, as that is the only case directly relevant for the follow-up discussion.

Field decomposition of the Dirac field

The free particle portion of the QED Lagrangian density for the Dirac field is

$$\mathcal{L} = i\bar{\psi}(\partial_\mu \gamma^\mu - m)\psi$$

which leads to the free Dirac Hamiltonian

$$H = \int \mathrm{d}^3\boldsymbol{x}\ \bar{\psi}(\boldsymbol{x})(-i\boldsymbol{\gamma}\cdot\nabla + m)\psi(\boldsymbol{x}) = \int \mathrm{d}^3\boldsymbol{x}\ \psi^\dagger(\boldsymbol{x})(-i\boldsymbol{\alpha}\cdot\nabla + \beta m)\psi(\boldsymbol{x}) \tag{2.1}$$

where using the form involving the matrices $\boldsymbol{\alpha} = \gamma^0\boldsymbol{\gamma}$ and $\beta = \gamma^0$ is common when working in the Hamiltonian formalism. As detailed in Appendix B, solving the differential equation corresponding to this Hamiltonian leads the four sets of solutions (for each three-momentum $\boldsymbol{p}$) — two with positive energy ($E \geq mc^2$) and two with negative energy ($E \leq mc^2$).

This set of solutions forms a complete basis of plane wave states and, naively[11], each one can be associated with a annihilation operator, with $\tilde{c}(r, \boldsymbol{p})$ for the positive energy states, $\tilde{d}(r, \boldsymbol{p})$ for the negative energy states, and where $r = 1, 2$ runs over the two orthogonal solutions for each

[11] As will be discussed in a few short paragraphs, it will be necessary to redefine these creation and annihilation operators to get the more physical field decomposition of the Dirac field. For that reason, the tildes distinguish these coefficients from the real creation and annihilation operators that will come later.

energy. Using the functions (B.4) and in appropriate normalization[12], the field operator can be decomposed as

$$\psi_a(\boldsymbol{x}) = \sum_r \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}} \frac{1}{\sqrt{2E_{\boldsymbol{p}}}} \left[\tilde{c}(r,\boldsymbol{p}) u_a^{(r)}(\boldsymbol{p}) + \tilde{d}(r,-\boldsymbol{p}) v_a^{(r)}(-\boldsymbol{p})\right] \mathrm{e}^{i\boldsymbol{p}\cdot\boldsymbol{x}}$$

where $E_{\boldsymbol{p}} = \sqrt{\boldsymbol{p}^2 + m^2}$ is the symbol for the positive energy corresponding to the three-momentum and mass, and all the while keeping in mind that it is a four-component function, i.e. $a = 1, 2, 3, 4$. It is the negative energy eigenvalue of the negative energy solutions that necessitates the minus signs for the momenta in this form, and alternatively the minus sign could also go into the exponent (i.e. $\mathrm{e}^{-i\boldsymbol{p}\cdot\boldsymbol{x}}$).

As electrons are fermions, the creation and annihilation operators obey the anti-commutation relations

$$\{\tilde{c}(\boldsymbol{p},r), \tilde{c}^\dagger(\boldsymbol{q},s)\} = \delta_{rs}\delta(\boldsymbol{p}-\boldsymbol{q}) = \{\tilde{d}(\boldsymbol{p},r), \tilde{d}^\dagger(\boldsymbol{q},s)\}$$

$$\{\tilde{c}(\boldsymbol{p},r), \tilde{d}^\dagger(\boldsymbol{q},s)\} = 0 = \{\tilde{d}(\boldsymbol{p},r), \tilde{c}^\dagger(\boldsymbol{q},s)\}$$

which can be shown would lead to the anti-commutation relation

$$\{\psi_a(\boldsymbol{x}), \psi_b^\dagger(\boldsymbol{y})\} = \delta_{ab}\delta(\boldsymbol{x}-\boldsymbol{y})$$

for the field operator. In this interpretation the operators $\tilde{c}(\boldsymbol{p},r)$ and $\tilde{d}(\boldsymbol{p},r)$ should annihilate the corresponding particle from a state, and also destroy a vacuum state $|\tilde{0}\rangle$, i.e.

$$\tilde{c}(\boldsymbol{p},r)\,|\tilde{0}\rangle = 0, \quad \tilde{c}(\boldsymbol{p},r)\,|\tilde{0}\rangle = 0$$

The operators $\tilde{c}^\dagger(\boldsymbol{p},r)$ and $\tilde{d}^\dagger(\boldsymbol{p},r)$ would add the corresponding particle into a state.

As the states we use to expand the field operator are the eigenstates of the Hamiltonian (2.1), it is possible to write down the Hamiltonian in terms of the creation and annihilation operators as

$$H = \sum_r \int \mathrm{d}\boldsymbol{p}\; E(\boldsymbol{p}) \left[\tilde{c}^\dagger(\boldsymbol{p},r)\tilde{c}(\boldsymbol{p},r) - \tilde{d}^\dagger(\boldsymbol{p},r)\tilde{d}(\boldsymbol{p},r)\right] \tag{2.2}$$

where the symbol $E(\boldsymbol{p}) = +\sqrt{\boldsymbol{p}^2 + m^2}$ corresponds is the *positive* energy corresponding to the three-momentum $\boldsymbol{p}$, and the negative energy states have appropriately received a minus sign. This Hamiltonian, however, has an obvious physical problem: if we keep adding particles that correspond to the $\tilde{d}^\dagger(\boldsymbol{p},r)$ operators, the energy of the state will become lower and lower. While these are negative energy states, *physically*, the energy should always increase when particles get added.

To resolve this conundrum, it is necessary to reinterpret the creation-annihilation operators and the vacuum state. As the particles obey anti-commutation relations and therefore only up to just one particle is allowed in any given state, there is not much difference (mathematically) between the creation and annihilation operators. It is therefore possible to flip the annihilation operator for the negative energy states into a creation operator (and vice versa for the creation operator) by defining

$$c(\boldsymbol{p},r) = \tilde{c}(\boldsymbol{p},r), \quad d(\boldsymbol{p},r) = \tilde{d}^\dagger(\boldsymbol{p},r) \tag{2.3}$$

[12]The normalization factor $1/\sqrt{2E_{\boldsymbol{p}}}$ comes from the requirement for the three-dimensional integrals to be Lorentz invariant. This can also be linked to four-dimensional space-time integrals without any explicit normalization factors, but which contain $\delta(p^2 - m^2)$ to constrain the four-momenta to the mass shell (see Section A.2 in the Appendix).

While clearly nothing changes for $c(\boldsymbol{p}, r)$, the anti-commutation relation for the $d(\boldsymbol{p}, r)$, due to the anti-commutator being symmetric, still holds

$$\{d(\boldsymbol{p}, r), d^\dagger(\boldsymbol{q}, s)\} = \{d^\dagger(\boldsymbol{p}, r), d(\boldsymbol{q}, s)\} = \{\tilde{d}(\boldsymbol{p}, r), \tilde{d}^\dagger(\boldsymbol{q}, s)\} = \delta_{rs}\delta(\boldsymbol{p} - \boldsymbol{q})$$

However, what changes is that the new set of creation and annihilation operators define a new vacuum state $|0\rangle$

$$c(\boldsymbol{p}, r)\,|0\rangle = 0, \quad d(\boldsymbol{p}, r)\,|0\rangle = 0$$

that corresponds to the physical vacuum with zero particles and zero energy.

To transform the Hamiltonian to the new operators, one can just plug the redefined operators (2.3) into the Hamiltonian (2.2). Using the anti-commutation relations, it is possible to get rid of the minus sign in the second part of the Hamiltonian

$$\begin{aligned} -\tilde{d}^\dagger(\boldsymbol{q}, s)\tilde{d}(\boldsymbol{p}, r) = -d(\boldsymbol{q}, s)d^\dagger(\boldsymbol{p}, r) = d^\dagger(\boldsymbol{p}, r)d(\boldsymbol{q}, s) + \delta_{rr}\delta(\boldsymbol{p} - \boldsymbol{p}) \\ = d^\dagger(\boldsymbol{p}, r)d(\boldsymbol{q}, s) + \infty \end{aligned}$$

Now, while the infinity is weird, it is not actually a problem — it is just a constant and we can shift the total energy by any constant without changing any of the physics. If we now just remove the infinity, the Hamiltonian becomes

$$\hat{H} = \sum_r \int \mathrm{d}\boldsymbol{p}\; E(\boldsymbol{p}) \left[c^\dagger(\boldsymbol{p}, r)c(\boldsymbol{p}, r) + d^\dagger(\boldsymbol{p}, r)d(\boldsymbol{p}, r)\right] \tag{2.4}$$

The key property of this form of the Hamiltonian is that the creation of any particle will always lead to an increase in energy, which makes much more sense physically.

Finally, for completeness, the field decomposition of the Dirac field in terms of the new creation and annihilation operators is

$$\psi_a(\boldsymbol{x}) \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}} \frac{1}{\sqrt{2E_{\boldsymbol{p}}}} \left[c(r, \boldsymbol{p})u_a^{(r)}(\boldsymbol{p}) + d^\dagger(r, -\boldsymbol{p})v_a^{(r)}(-\boldsymbol{p})\right] \mathrm{e}^{i\boldsymbol{p}\cdot\boldsymbol{x}} \tag{2.5}$$

Both $c^\dagger$ and $d^\dagger$ create particles, but they create different particles. After the reinterpretation, it is much easier to accept that the particles associated with the negative energy solutions of the free-particle Dirac equation are in fact simply positrons. The reinterpreted operator $d^\dagger$ creates particles that have positive energy, which is what one would expect from a massive particle with kinetic energy. The field creation operator $\psi^\dagger(\boldsymbol{x})$, therefore, creates an electron but *annihilates a positron*. Similarly, the operator $\psi(\boldsymbol{x})$ annihilates an electron and creates a positron.

However, what was wrong with the initial operators $\tilde{d}$ and $\tilde{d}^\dagger$, the and Hamiltonian (2.2)? Actually, technically nothing. They simply do not have the physical interpretation one would naively expect. Specifically, the positron operator $\tilde{d}(\boldsymbol{p}, r) = d^\dagger(\boldsymbol{p}, r)$ was not a positron annihilation operator, but a positron creation operator. But this, in turn, means that the original vacuum state $|\tilde{0}\rangle$ was already fully filled up with positrons — the so-called *Dirac sea* — as adding any additional (physical) positrons would annihilate it. And so the dynamics predicted by the negative energy solutions of the Dirac equation are, in fact, the dynamics of the absence of the corresponding positron states (holes) in the so-called infinitely charged Dirac sea of positrons, and not of physical positrons or electrons.

Particle number and charge operators

Particularly useful operators that can be defined with the help of the creation and annihilation operators are the particle number operators

$$N_{\mathrm{e}} = \sum_r \int \mathrm{d}^3\boldsymbol{p}\; c^\dagger(\boldsymbol{p}, r)c(\boldsymbol{p}, r), \quad N_{\mathrm{p}} = \sum_r \int \mathrm{d}^3\boldsymbol{p}\; d^\dagger(\boldsymbol{p}, r)d(\boldsymbol{p}, r)$$

where N_{e} counts the number of electrons and N_{p} counts the number of positrons in a given state. Of course, the sum of those operators

$$N = N_{\mathrm{e}} + N_{\mathrm{p}} = \sum_r \int \mathrm{d}^3\boldsymbol{p}\;\left[c^\dagger(\boldsymbol{p}, r)c(\boldsymbol{p}, r) + d^\dagger(\boldsymbol{p}, r)d(\boldsymbol{p}, r)\right]$$

counts the total number of particles in a state.

However, as will be shown in Section 2.3, the electron-positron pair creation and annihilation effect of the electromagnetic interaction means that the total number of particles is not a good quantum number in QED. Instead, it is convenient to consider the conserved charged related to the $U(1)$ gauge symmetry of the Dirac field (and the elecromagnetic interaction), which from the gauge transformation

$$\psi \to \mathrm{e}^{i\alpha}\psi \tag{2.6}$$

leads to a conserved charge four-current density

$$j^\mu = \bar{\psi}\gamma^\mu\psi$$

Integrating the time-like component gives the total electric charge of the system

$$\begin{aligned} Q = J^0 = \int \mathrm{d}^3\boldsymbol{x}\; j^0(\boldsymbol{x}) &= \int \mathrm{d}^3\boldsymbol{x}\; \psi^\dagger(\boldsymbol{x})\psi(\boldsymbol{x}) \\ &= \sum_r \int \mathrm{d}^3\boldsymbol{p}\;\left[c^\dagger(r, \boldsymbol{p})c(r, \boldsymbol{p}) + d(r, -\boldsymbol{p})d^\dagger(r, -\boldsymbol{p})\right] \\ &\simeq \sum_r \int \mathrm{d}^3\boldsymbol{p}\;\left[c^\dagger(r, \boldsymbol{p})c(r, \boldsymbol{p}) - d^\dagger(r, \boldsymbol{p})d(r, \boldsymbol{p})\right] \end{aligned} \tag{2.7}$$

where in the last step commuting the operators also leads to an infinity which has been dropped from the expression. This is a conserved quantity in QED, and, crucially, it is clear from the minus sign that electrons and positrons have the opposite charge.

2.3 Electromagnetic interaction in QED

The previous section did not address the interactions between the fields. The standard and general way of introducing an interaction to a gauge field theory is to perform the *minimal substitution*, that is replacing the derivatives with the *gauge covariant derivative*[13]. In QED, the covariant derivative introducing the electromagnetic field A^μ is defined as

$$\partial^\mu \to D^\mu = \partial^\mu + ieA^\mu$$

[13] "Covariance" here refers to the being covariant with respect to the gauge transformation under a symmetry group. In the case of the electromagnetic interactions, this is the simple $U(1)$ group transforming according to (2.6), but this generalises to other interactions described by more complex symmetry groups, such as the $SU(2)$ for the weak interaction and $SU(3)$ for the strong interaction.

and leads to the following interaction terms

$$\mathcal{L}_{\text{int}} = -eA_\mu \bar{\psi}\gamma^\mu\psi, \quad H_{\text{int}} = e\int \mathrm{d}^3\boldsymbol{x}\ \bar{\psi}(\boldsymbol{x})\gamma^\mu\psi(\boldsymbol{x})A_\mu(\boldsymbol{x}) \tag{2.8}$$

in the Lagrangian and the Hamiltonian. The elementary charge e acts as the interaction coupling constant[14], determining the strength of the interaction. If we split the elecromagnetic field up into its time- and space-like components $A^\mu = (\phi, \boldsymbol{A})$ the corresponding interaction term in Hamiltonian[15] becomes

$$H_{\text{int}} = e\int \mathrm{d}^3\boldsymbol{x}\ \bar{\psi}(\boldsymbol{x})\gamma^0\psi(\boldsymbol{x})\phi(\boldsymbol{x}) - e\int \mathrm{d}^3\boldsymbol{x}\ \bar{\psi}(\boldsymbol{x})\boldsymbol{\gamma}\cdot\boldsymbol{A}(\boldsymbol{x})\psi(\boldsymbol{x})$$

Interpretation of the interaction term

By plugging the field decomposition of the Dirac field (2.5) into interaction Hamiltonian (2.8), it is possible to get a better physical sense of the effect the interaction term. In order to interpret it, it is useful to remember that the Hamiltonian operator generates the time-evolution of a state. In other words, given an initial state $|\Psi(0)\rangle$, at time t the state $|\Psi(t)\rangle$ can be calculated by

$$|\Psi(t)\rangle = \exp\left(-iHt\right)|\Psi(0)\rangle = \left[1 - iHt + \mathrm{O}(t^2)\right]|\Psi(0)\rangle$$

where in the second part the exponential has been Taylor-expanded[16]. The expansion allows reasoning about what happens to the initial state at very small (or even infinitesimal) time scales. Initially, the state stays the same as it was at $t = 0$ due to the constant unity. However, the $-iHt$ part introduces a tiny rotation in the state generated by the Hamiltonian H, with its magnitude proportional to the time step t. Because of this, it is sufficient to analyse the effect of each term in the Hamiltonian H separately and still get a qualitative understanding of how the system behaves.

In order to analyse the interaction term in the Dirac Hamiltonian, it is convenient to split the Dirac field decomposition into two parts

$$\psi_+(\boldsymbol{x}) = \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}}\frac{1}{\sqrt{2E_{\boldsymbol{p}}}}c(r,\boldsymbol{p})u^{(r)}(\boldsymbol{p})\mathrm{e}^{i\boldsymbol{p}\cdot\boldsymbol{x}} \tag{2.9}$$

$$\psi_-(\boldsymbol{x}) = \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}}\frac{1}{\sqrt{2E_{\boldsymbol{p}}}}d^\dagger(r,\boldsymbol{p})v^{(r)}(\boldsymbol{p})\mathrm{e}^{-i\boldsymbol{p}\cdot\boldsymbol{x}} \tag{2.10}$$

such that we can write the whole field down as

$$\psi(x) = \psi_+(\boldsymbol{x}) + \psi_-(\boldsymbol{x})$$

[14] The sign in the covariant derivative differs between different sources, and depends on (1) the sign convention in the Minkowski metric, (2) the sign convention of the electric charge, and (3) the sign convention in the exponential of the gauge transformation. In this thesis, the elementary charge e is assumed to be a positive value, but by convention the charge of the electron is negative $q_e = -e$. This, together with the $(+,-,-,-)$ metric signature and the sign choice in (2.6), leads to the + sign in the covariant derivative. However, in other sources the gauge covariant derivative may be defined as $D^\mu = \partial^\mu - iqA^\mu$.

[15] Cf. pg. 123, eq. 4.129 in Peskin and Schroeder [24].

[16] Actually, it would be more correct to say that we have replaced the exponential with its *definition* in terms of an infinite sum — functions of operators are generally defined by the Taylor series of their real or complex function counterparts.

This is convenient as $\psi_+(\boldsymbol{x})$ purely annihilates electrons, and $\psi_-(\boldsymbol{x})$ purely creates positrons. Of course, correspondigly $\psi_+^\dagger(\boldsymbol{x})$ and $\psi_-^\dagger(\boldsymbol{x})$ create particles and annihilate antiparticles, respectively. Keeping the $\boldsymbol{x}$-dependence of all the fields (ψ, $\psi_\pm$, A_μ) implicit for space reasons, the interaction in the Hamiltonian can be expanded as

$$H_{\text{int}} = e\int \mathrm{d}^3\boldsymbol{x}\ \bar{\psi}\gamma^\mu\psi A_\mu = e\int \mathrm{d}^3\boldsymbol{x}\ [\bar{\psi}_+\gamma^\mu\psi_+ + \bar{\psi}_+\gamma^\mu\psi_- + \bar{\psi}_-\gamma^\mu\psi_+ + \bar{\psi}_-\gamma^\mu\psi_-]A_\mu \tag{2.11}$$

leading to four distinct terms with slightly different effects.

The $\boldsymbol{x}$-dependent products can be expanded and simplified further using the field decompositions (2.9) and (2.10) of the Dirac field, allowing for more precise interpretation. The terms $\bar{\psi}_+\gamma^\mu\psi_+$ and $\bar{\psi}_-\gamma^\mu\psi_-$ represent the standard effect that an electromagnetic field has on a charged particle: it nudges them in some way. This can be seen from

$$\bar{\psi}_+\gamma^\mu\psi_+ = \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}}\int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}}\sum_{r,r'}\frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}}\bar{u}^{(r)}(\boldsymbol{p})\gamma^\mu u^{(r')}(\boldsymbol{p}')\mathrm{e}^{i(\boldsymbol{p}-\boldsymbol{p}')\cdot\boldsymbol{x}}c^\dagger(r,\boldsymbol{p})c(r',\boldsymbol{p}')$$

where the creation-annihilation structure $c_k^\dagger c_l$ has the effect of removing a momentum state and replacing it with another one. How the elecromagnetic field exactly influences the electron state depends on the momentum-composition of the electromagntic field. This can be seen by adding on the electromagnetic field and integrating over space, and swapping the order of integration

$$\begin{aligned} e\int \mathrm{d}^3\boldsymbol{x}\ \bar{\psi}_+\gamma^\mu\psi_+A_\mu = e\int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}}\int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}}\sum_{r,r'}\frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}}\bar{u}^{(r)}(\boldsymbol{p})\gamma^\mu u^{(r')}(\boldsymbol{p}') \\ \cdot\ c^\dagger(r,\boldsymbol{p})c(r',\boldsymbol{p}')\int \mathrm{d}^3\boldsymbol{x}\ A_\mu(\boldsymbol{x})\mathrm{e}^{-i(\boldsymbol{p}-\boldsymbol{p}')\cdot\boldsymbol{x}} \end{aligned}$$

The $\int \mathrm{d}^3\boldsymbol{x}\ A_\mu(\boldsymbol{x})\mathrm{e}^{-i(\boldsymbol{p}-\boldsymbol{p}')\cdot\boldsymbol{x}}$ part is a Fourier transform over space of the electromagnetic field A_μ, giving us the magnitude of each momentum component of the electromagnetic field. Similarly, we can calculate term which nudges the antiparticles

$$\begin{aligned} \bar{\psi}_-\gamma^\mu\psi_- &= \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}}\int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}}\sum_{r,r'}\frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}}\bar{v}^{(r)}(\boldsymbol{p})\gamma^\mu v^{(r')}(\boldsymbol{p}')\mathrm{e}^{-i(\boldsymbol{p}-\boldsymbol{p}')\cdot\boldsymbol{x}}d(r,\boldsymbol{p})d^\dagger(r',\boldsymbol{p}') \\ &= \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}}\int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}}\sum_{r,r'}\frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}}\bar{v}^{(r)}(\boldsymbol{p})\gamma^\mu v^{(r')}(\boldsymbol{p}')\mathrm{e}^{-i(\boldsymbol{p}-\boldsymbol{p}')\cdot\boldsymbol{x}}\delta_{r,r'}\delta^{(3)}(\boldsymbol{p}-\boldsymbol{p}') \\ &\quad - \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}}\int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}}\sum_{r,r'}\frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}}\bar{v}^{(r)}(\boldsymbol{p})\gamma^\mu v^{(r')}(\boldsymbol{p}')\mathrm{e}^{-i(\boldsymbol{p}-\boldsymbol{p}')\cdot\boldsymbol{x}}d^\dagger(r',\boldsymbol{p}')d(r,\boldsymbol{p}) \\ &= \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}}\sum_{r}\frac{1}{(2\pi)^3 2E_{\boldsymbol{p}}}\bar{v}^{(r)}(\boldsymbol{p})\gamma^\mu v^{(r)}(\boldsymbol{p}) \\ &\quad - \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}}\int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}}\sum_{r,r'}\frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}}\bar{v}^{(r)}(\boldsymbol{p})\gamma^\mu v^{(r')}(\boldsymbol{p}')\mathrm{e}^{-i(\boldsymbol{p}-\boldsymbol{p}')\cdot\boldsymbol{x}}d^\dagger(r',\boldsymbol{p}')d(r,\boldsymbol{p}) \end{aligned}$$

except that commuting the creation-annihilation operators in order to achieve normal ordering leads to a minus sign, indicating that the positrons are affected in the opposite way compared to the electrons, consistent with their opposite charge. The other integral, arising from the delta function, is just another constant infinity that has no effect on the physics and can be ignored here.

The other two terms in (2.11) are physically much more interesting. Writing them out in terms of the creation and annihilation operators

$$\bar{\psi}_+\gamma^\mu\psi_- = \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}} \int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}} \sum_{r,r'} \frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}} \bar{u}^{(r)}(\boldsymbol{p})\gamma^\mu v^{(r')}(\boldsymbol{p}')\mathrm{e}^{i(\boldsymbol{p}+\boldsymbol{p}')\cdot\boldsymbol{x}} c^\dagger(r,\boldsymbol{p})d^\dagger(r',\boldsymbol{p}')$$

$$\begin{aligned}\bar{\psi}_-\gamma^\mu\psi_+ &= \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}} \int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}} \sum_{r,r'} \frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}} \bar{v}^{(r)}(\boldsymbol{p})\gamma^\mu u^{(r')}(\boldsymbol{p}')\mathrm{e}^{-i(\boldsymbol{p}+\boldsymbol{p}')\cdot\boldsymbol{x}} d(r,\boldsymbol{p})c(r',\boldsymbol{p}') \\ &= -\int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}} \int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}} \sum_{r,r'} \frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}} \bar{v}^{(r)}(\boldsymbol{p})\gamma^\mu u^{(r')}(\boldsymbol{p}')\mathrm{e}^{-i(\boldsymbol{p}+\boldsymbol{p}')\cdot\boldsymbol{x}} c(r',\boldsymbol{p}')d(r,\boldsymbol{p})\end{aligned}$$

show that the terms have the effect of creating ($c_k^\dagger d_l^\dagger$) or annihilating ($c_k d_l$) an electron-positron pair in the state. In other words, in the presence of an electromagnetic field, an initial state with a known number of positrons and electrons (e.g. made up purely of a fixed number of electrons) will, over time, accumulate trace amounts of components in the wavefunction with different particle numbers.

In particular, given an initial (N, M)-particle state of N electrons and M positrons, after an infinitesimal amount of time it will contain components with $(N+1, M+1)$ and $(N-1, M-1)$ particles as well. As time goes on, those components can then connect to subspaces with other particle numbers, as long as it is some number of electron-positron pairs away (e.g. $(N+2, M+2)$ or $(N-3, M-3)$ particles). This is consistent with the conservation of charge mentioned at the end of Section 2.2 — an electron-positron pair has zero total charge, and so as long as only pairs are created or annihilated, the total charge of the system remains unchanged. However, what is clear is that the total number of particles in a state is no longer a good quantum number.

Free Dirac field term. The preceding discussion focused on the interaction and did not touch upon the free particle part of the Hamiltonian

$$\hat{H}_0 = \sum_r \int \mathrm{d}\boldsymbol{p}\ E(\boldsymbol{p}) \left[c^\dagger(\boldsymbol{p},r)c(\boldsymbol{p},r) + d^\dagger(\boldsymbol{p},r)d(\boldsymbol{p},r)\right]$$

In the plane wave representation like this, which clearly separates the electron and positron parts, the effect of this term on a state is quite simple: each plane wave component gets a phase shift proportional to its energy. Most crucially, however, this part of the term introduces no mixing between the electron and positron parts of the wavefunction.

Interaction term and the photon field

It should not be forgotten that that in QED $A^\mu(x)$ is also a dynamical quantum field associated with particles, the photons. As such, just as a non-zero photon field can affect the electron-positron part of the wavefunction, so can a non-zero Dirac field affect the photon part. To clearly see this, it is convenient to pack up all the Dirac parts of the Hamiltonian into the charge four-current of the Dirac field

$$j^\mu(x) = \bar{\psi}(x)\gamma^\mu\psi(x)$$

and then rewrite the interaction term in the Hamiltonian (2.8) as

$$\begin{aligned} H_{\text{int}} &= e \int \mathrm{d}^3\boldsymbol{x}\ j_\mu(t,\boldsymbol{x}) A^\mu(t,\boldsymbol{x}) \\ &= \int \mathrm{d}^3\boldsymbol{p} \sum_{r=1}^{3} \frac{1}{\sqrt{E_{\boldsymbol{p}}}} \left(a_{r,-\boldsymbol{p}} \epsilon_r^\mu(-p) \mathrm{e}^{-i|p|t} + a_{r,\boldsymbol{p}}^\dagger \epsilon_r^{\mu *}(p) \mathrm{e}^{i|p|t} \right) \int \mathrm{d}^3\boldsymbol{x}\ j_\mu(t,\boldsymbol{x}) \mathrm{e}^{-i\boldsymbol{p}\cdot\boldsymbol{x}} \end{aligned}$$

where in the second part the photon field was expanded with the help of the field decomposition of A_μ [24, Eq. (4.131)]. As can be seen, the mere presence of a non-zero charge density from the Dirac field will cause photons to be created an annihilated in the wavefunction. In fact, if one were to expand j^μ as well, it would be clear that there are two distinct ways of creating or annihilating photons: (1) by changing the momentum of electrons and positrons (or, more generally, any electrically charged particle), and (2) by the creating or annhilating particle-antiparticle pairs.

2.4 Classical external field and stationary states

One useful modification (or approximation, depending on the circumstances) is the introduction of a classical external electromagnetic field $A^\mu_{\text{ext}}(x)$ to the equations. This can be achieved by substituting

$$A^\mu(x) \to A^\mu(x) + A^\mu_{\text{ext}}(x) \tag{2.12}$$

and keeping in mind that $A^\mu(x)$ is still a fully quantized field, whereas $A^\mu_{\text{ext}}(x)$ is an externally defined (time-dependent) value. Crucially, while it will affect the dynamics of the electrons, it itself will not be affected by the behaviour of the charged particles. This procedure is useful because it is often not necessary to concern ourselves with all the quantum mechanical details of the photon field. Practical examples of this procedure include replacing the atomic nuclei with simple electrostatic potentials and the treatment of external laser fields in ionization studies.

It can also be useful to think of the external field as being generated by an external four-current density $j^\mu_{\text{ext}}(x)$. This, for example, justifies the use of a potential derived from a nuclear charge distribution. In this case, it is assumed that the external classical component of the electromagnetic field is related to the charge four-current via Maxwell's equations

$$\partial_\nu \partial^\nu A^\mu_{\text{ext}}(x) = \mu_0 j^\mu_{\text{ext}}(x)$$

assumed to be in Lorenz gauge $\partial_\mu A^\mu_{\text{ext}} = 0$. In fact, instead of ever introducing an external electromagnetic field, it is possible to incorporate the external field into QED by substituting the charge four-current

$$j^\mu(x) \to j^\mu(x) + j^\mu_{\text{ext}}(x)$$

of the Dirac field. This would lead to an interaction term between the quantized photons and the external charge four-current

$$H_{\text{int,ext}} = e \int \mathrm{d}^3\boldsymbol{x}\ j^\mu_{\text{ext}}(t,\boldsymbol{x}) A_\mu(\boldsymbol{x})$$

Interestingly, while there is no interaction between photons and the external charge when substituting according to (2.12), these two approaches are actually physically equivalent [49, Chapter 14][41].

External field Dirac equation

Keeping the classical external electromagnetic field but dropping the quantized field leads to the external field Dirac equation. Formally, this can be done simply by removing the photon field from the QED Lagrangian, arriving at[17]

$$\mathcal{L} = i\bar{\psi}(\gamma^\mu \partial_\mu - m)\psi - e\bar{\psi}\gamma^\mu \psi A^{\text{ext}}_\mu$$

Further separating the time-like scalar potential and the space-like vector potential of the electromagnetic potential four-vector $A^\mu_{\text{ext}}(t, \boldsymbol{x}) = (\phi_{\text{ext}}(t, \boldsymbol{x}), \boldsymbol{A}_{\text{ext}}(t, \boldsymbol{x}))$ and applying a Legendre transformation leads to the more familiar Hamiltonian form

$$H = \int \mathrm{d}^3\boldsymbol{x}\ \bar{\psi}\left[-i\boldsymbol{\gamma}\cdot\nabla + m + e\gamma^0\phi_{\text{ext}}(t,\boldsymbol{x}) - e\boldsymbol{\gamma}\cdot\boldsymbol{A}_{\text{ext}}(t,\boldsymbol{x})\right]\psi \tag{2.13}$$

The part in the brackets, when γ^0 is moved out of $\bar{\psi}$, is the single-particle external field Dirac equation

$$i\frac{\partial\psi}{\partial t} = \hat{h}_{\boldsymbol{D},\text{ext}}\psi, \quad \hat{h}_{\boldsymbol{D},\text{ext}} = -i\boldsymbol{\alpha}\cdot\nabla + e\phi_{\text{ext}}(t,\boldsymbol{x}) - e\boldsymbol{\alpha}\cdot\boldsymbol{A}_{\text{ext}}(t,\boldsymbol{x}) + \beta m$$

where the gamma matrices have been replaced by $\boldsymbol{\alpha} = \gamma^0\boldsymbol{\gamma}$ and $\beta = \gamma^0$, which is the more conventional representation in the Hamiltonian form, especially in electronic structure theory.

A particular special case of this equation relevant for this thesis is when the external field is simply a time-independent potential $A^\mu_{\text{ext}}(t, \boldsymbol{x}) = (V(\boldsymbol{x}), 0)$, such as the one generated by some stationary charge distribution, such as an atomic nucleus. This leads to the Hamiltonian

$$\hat{h}_{\boldsymbol{D},\text{ext}} = -i\boldsymbol{\alpha}\cdot\nabla + \beta m + eV(\boldsymbol{x}) \tag{2.14}$$

The single-particle Hamiltonian operator (2.14) has no explicit time-dependence and so its eigenstates are the stationary states, approximating the states of a single-electron system. Assuming a reasonable, unconfined[18] potential $V(\boldsymbol{x})$, such as one generated by a nuclear charge distribution (see also Section 2.7), the spectra of the Hamiltonian generally has three sets of states: (1) a positive continuum with energies $E \geq mc^2$, (2) a negative continuum with energies $E \leq -mc^2$, and (3) a discrete set of bound states, with energies $-mc^2 < E < mc^2$. These bound states correspond to the different energy levels an electron captured by the atom can be on.

Interpretation of the stationary states

There is an interpretational conundrum that rarely seems to be discussed in literature or textbooks. If QED and its pair creation-annihilation aspects are ignored, and just the basic single-particle Dirac Hamiltonian of the form (2.14) is considered, it is generally assumed to represent a theory where every state is an electron.

However, as was discussed in Section 2.3, an interaction term with the electromagnetic field leads to pair creation and annihilation terms in the Hamiltonian. This also holds for the interaction with a classical external electromagnetic field A^{ext}_μ, which is clear as the interaction term

$$H^{\text{ext}}_{\text{int}} = e\int \mathrm{d}^3\boldsymbol{x}\ \bar{\psi}\gamma^\mu\psi A^{\text{ext}}_\mu = e\int \mathrm{d}^3\boldsymbol{x}\ [\bar{\psi}_+\gamma^\mu\psi_+ + \bar{\psi}_+\gamma^\mu\psi_- + \bar{\psi}_-\gamma^\mu\psi_+ + \bar{\psi}_-\gamma^\mu\psi_-]A^{\text{ext}}_\mu \tag{2.15}$$

[17] In order to make the equations more readable, the $\square_{\text{ext}}$ label is liberally moved up or down depending on where the four-vector index needs to go.

[18] It is interesting to note that confinement potentials lead to problems with the Dirac equation if introduced in the potential term instead as a scalar potential added to the mass term [50].

is essentially identical to (2.11) and these matrix elements are indeed taken into account when (2.14) is diagonalized. The implication is that an external field should therefore also couple subspaces with different particle numbers, which are some number of electron-positron pairs away. In fact, it should couple to subspaces with ever-increasing number of particles.

Herein lies the problem: the differential operator (2.14) should be equivalent to the many-particle Hamiltonian (2.13) in the $(1,0)$-particle case, i.e. with one electron and no positrons. However, it is quite clear that can not be the case if the many-particle Hamiltonian is allowed to couple spaces with different particle numbers together. A way to think about it is to consider the number of degrees of freedom a many-particle state has when expanded in terms of all the many-particle basis product states, and compare that to the expansion of a single-electron wavefunction in terms of the single-particle states — the number of degrees of freedom in the former case is much greater.

It is possible to reconcile this problem though. When the Dirac equation is generalized to a many-particle QFT, there was one step that was not directly justified from the Euler-Lagrange equations — the redefinition (2.3) of the creation and annihilation operators. This was to fix the problem of negative energy solutions of the free-particle Dirac Hamiltonian, which do not make sense physically, changing their interpretation. And as a consequence, it also lead to the redefinition of the vacuum state. But there was actually mathematically nothing wrong with the original choice of the creation and annihilation operators, and the redefinitions was largely just an exercise in choosing a new symbol.

The original operators $\tilde{c}(r,\boldsymbol{p})$ and $\tilde{d}(r,\boldsymbol{p})$ correspond to the positive and negative energy states of the Dirac equation, respectively. If we write the pair creation and pair annihilation parts of the Hamiltonian (2.15) in terms of the original operators

$$\bar{\psi}_+\gamma^\mu\psi_- = \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}} \int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}} \sum_{r,r'} \frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}} \bar{u}^{(r)}(\boldsymbol{p})\gamma^\mu v^{(r')}(\boldsymbol{p}')\mathrm{e}^{i(\boldsymbol{p}+\boldsymbol{p}')\cdot\boldsymbol{x}}\tilde{c}^\dagger(r,\boldsymbol{p})\tilde{d}(r',\boldsymbol{p}')$$

$$\bar{\psi}_-\gamma^\mu\psi_+ = \int \frac{\mathrm{d}^3\boldsymbol{p}}{\sqrt{(2\pi)^3}} \int \frac{\mathrm{d}^3\boldsymbol{p}'}{\sqrt{(2\pi)^3}} \sum_{r,r'} \frac{1}{\sqrt{2E_{\boldsymbol{p}}}\sqrt{2E_{\boldsymbol{p}'}}} \bar{v}^{(r)}(\boldsymbol{p})\gamma^\mu u^{(r')}(\boldsymbol{p}')\mathrm{e}^{-i(\boldsymbol{p}+\boldsymbol{p}')\cdot\boldsymbol{x}}\tilde{d}^\dagger(r,\boldsymbol{p})\tilde{c}(r',\boldsymbol{p}')$$

the terms no longer appear to mix particle particle number. Instead, it is clear that these terms simply mix the positive and negative energy states together. Nothing very interesting happens to the pure electron and pure positron terms $\bar{\psi}_+\gamma^\mu\psi_+$ and $\bar{\psi}_-\gamma^\mu\psi_-$ — they simply keep nudging their corresponding particles. Although it is also possible to interpret the latter as nudging the negative energy states.

It is possible to give a physical interpretation to the original formulation that can explain the apparent inconsistency of the pair creation and annihilation terms with the Dirac equation. For this it is necessary to remember that the "vacuum" state $|\tilde{0}\rangle$ corresponding to the operators $\tilde{c}(r,\boldsymbol{p})$ and $\tilde{d}(r,\boldsymbol{p})$ is in fact physically (i.e. in the proper QED interpretation) non-empty and filled with positrons — adding another one with $d^\dagger(r,\boldsymbol{p}) = \tilde{d}(r,\boldsymbol{p})$ would by definition annihilate the state. Adding a positive energy state, i.e. $\tilde{c}^\dagger(r,\boldsymbol{p})\,|\tilde{0}\rangle$, would create a state that physically has an infinity of positrons and one electron, or in other words a $(1,\infty)$-particle state. On the other hand, adding a negative energy state, i.e. $\tilde{d}^\dagger(r,\boldsymbol{p})\,|\tilde{0}\rangle$, annihilates one of the positrons and leads to a $(0,\infty-1)$-particle state[19]. To construct physically interesting states of only electrons (e.g. 1-electron/0-positron) it would be necessary to first annihilate all the positrons in the $|\tilde{0}\rangle$ "vacuum". But, in principle, it would be completely valid to generate the full Fock space this

[19] Alternatively, rather than counting the number of positrons down from infinity, it can be more convenient to talk about "holes". So this would be a "0-electron/1-positron-hole state".

way with the original naive creation and annihilation operators — the dynamics will be the same, even if the positron labelling would look weird.

With all that in mind, it is worth contemplating the interpretation of a single particle state

$$|\Psi\rangle = \sum_r \int \mathrm{d}^3\boldsymbol{p} \left[a(r,\boldsymbol{p})\tilde{c}^\dagger(r,\boldsymbol{p}) + b(r,\boldsymbol{p})\tilde{d}^\dagger(s,\boldsymbol{q})\right]|\tilde{0}\rangle \tag{2.16}$$

that is a superposition of both positive and negative energy states. The mixing between positive and negative energy states introduced by the external potential in the external field Dirac equation (2.14) means that, for example, the usual single-particle bound state solutions are exactly such superpositions. However, as these states are defined relative to the non-empty vacuum $|\tilde{0}\rangle$, such a states is in fact a superpositions of $(1,\infty)$-particle state (i.e. a single electron in a filled sea of positrons), and a $(0,\infty-1)$-particle state (annihilated positron hole in the filled sea).

Understanding the Dirac equation in this way also resolves the problem with the degrees of freedom. As the pair creation and pair annihilation terms only allow for coupling between particle number subspaces that are a "pair away", the $(1,\infty)$-particle (or positive energy) subspace can only connect to the $(0,\infty-1)$-particle subspace via pair annihilation, which is the negative energy subspace. Conversely, adding a pair to the negative energy subspace takes us into the positive energy subspace. Further annihilating a pair from the negative energy subspace is not possible as there are no more electrons to annihilate. Similarly, it is not possible to add a pair to the positive energy subspace because the positron component is already full. In other words, with $|\tilde{0}\rangle$ as the reference vacuum, the Hamiltonian is inherently restricted to what look like single-particle states in terms of degrees of freedom.

It also gives a physically consistent interpretation to the energies of the negative energy states and why they are more and more negative as the momentum increases. As those states, relative to the vacuum $|\tilde{0}\rangle$, correspond to the *removal* of a positron with some mass and kinetic energy, the energy must be lower than $-mc^2$. The higher the momentum (i.e. kinetic energy) of the state is, the more the energy is reduced. One thing that should be clarified is that this applies to the free-particle states — an interaction with the electromagnetic potential can lower the energies further.

An open question, however, is why the Dirac equation (2.14), while seemingly physically nonsensical, still yields very reasonable results. One likely reason is that, while the bound states do mix in the negative energy subspace, they do so relatively little (e.g. when measured as overlaps with the negative energy subspace), and the filled "positron sea" form a homogeneous background charge distribution which does not affect dynamics. Another thing to consider is that a pure negative energy state (a positron hole) is also negatively charged, which can be seen if the charge operator (2.7) is written in terms of $\tilde{c}(r,\boldsymbol{p})$ and $\tilde{d}(r,\boldsymbol{p})$

$$Q \simeq \sum_r \int \mathrm{d}^3\boldsymbol{p} \ \left[\tilde{c}^\dagger(r,\boldsymbol{p})\tilde{c}(r,\boldsymbol{p}) + \tilde{d}^\dagger(r,\boldsymbol{p})\tilde{d}(r,\boldsymbol{p})\right]$$

and so should resemble an electron dynamically.

2.5 Effective Hamiltonians and other approximations

The goal of theoretical atomic physics is to explain the interactions of the particles in atoms and to determine their electronic wavefunctions to predict atomic properties. Based on their theoretical treatment, there are broadly speaking two classes of problems that are solved in atomic physics:

1. Time-dependent processes, where we are looking at the evolution of the system over time. Physical examples include the scattering of particles or the ionization of electrons with laser pulses. Essentially, it is assumed that the system starts in some known configuration and then evolves in time, with theory used to predict how that evolution occurs.

2. Stationary states, where the Hamiltonian is diagonalized and the system is studied through its eigenstates, which are the stationary states of the system. The examples of this would be predicting energy levels and bound state wavefunctions.

There is also overlap between those two broad categories. For example, the stationary states of a time-dependent Hamiltonian at every time step are often a crucial part in time-dependent calculations. On the other hand, an argument can be made that every process in physics ultimately is a time-dependent process, with the eigenstates of the Hamiltonian just being a method of gaining insight into the underlying mathematical structure of the system.

However, performing some sort of "full QED" calculations in either of these categories is highly problematic. A naive numerical approach would involve iterating over combinations of single-particle states in some form, which is especially problematic for quantized photons. In large part due to their massless nature, the number of photons in a system is very fluid — each photon can be split up into more photons of less energy, and new photons would be constantly generated by the interaction term. When attempting to represent the precise photon wavefunction in some basis, the number of degrees of freedom would make the problem computationally intractable. [51, 52]

While not as pronounced, the fermionic matter sector of the theory suffers from similar issues. As electrons and positrons are massive, generating more and more particles comes with a severe energy penalty, but the pair-creation nature of the interaction term still means that basis determinants with higher and higher numbers of particles would have to be taken into account. Even if the particle number is fixed completely, as is usually done in most many-electron approaches, a good numerical description of a many-electron quantum state is difficult due to the combinatorial expansion of the basis and the entanglement, or correlation, between all the particles. However, generally speaking, it is possible to treat the fermionic wavefunction accurately, for low enough numbers of particles.

In summary, while solving the "full QED" problem is not necessarily feasible, there are various approaches and approximations that can lead to approximate but still precise and useful results for atomic systems.

Perturbation theory

A very common approach in QED is to employ perturbation theory, where the solutions being sought are written down as an infinite series, which ideally converges to the exact full QED answer. Each term of the series by itself is easier to calculate than the full solution, but the series usually has to be truncated somewhere, hoping that the partial sum leads to an accurate enough result, making the approach approximate.[20]

Most treatments of this topic present two types of perturbation theory:

1. Time-dependent perturbation theory which, as the name suggests, is used for time-dependent processes, such as the calculation of transition amplitudes.

2. Time-independent or stationary state perturbation theory, which is used to calculate the eigenvalues and eigenstates of a system.

[20] Most QED and QFT textbooks focus on perturbative approaches, especially the ones targeted at particle physicists. The main book used as a basis for this thesis was the one by Peskin and Schröder. [24]

This work primarily deals with stationary bound states, which implies that time-independent perturbation theory should be used. However, it turns out that in relativistic quantum field theory, there is a strong connection between the two formalisms.

General formalism

The common starting point for perturbation theory is to split the Hamiltonian[21] into two parts: the unperturbed part $\hat{H}_0$ and the perturbation $\hat{H}_\mathrm{I}$, such that the total Hamiltonian is written as

$$\hat{H} = \hat{H}_0 + \hat{H}_\mathrm{I}$$

Often these two parts correspond to the non-interacting free particles and interparticle interactions and are therefore called "free" and "interaction" parts, respectively,

The eigenstates and eigenvalues of $\hat{H}_0$ need to be determined using without perturbation theory, such as solving the differential equation analytically if possible. The solutions can be something as simple as the analytically known plane-wave solutions of free non-interacting particles, but can also be more complex, such as bound states with a non-standard potential, requiring a numerical approach. Once the solutions to the free part are known, the contribution of the interaction term to the full problem under study can be calculated term by term.

In fact, formally, it is convenient to introduce a parameter λ which can be used to turn the interaction on or off such that

$$\hat{H} = \hat{H}_0 + \lambda\hat{H}_\mathrm{I}$$

If $\lambda = 0$, the solution to the problem will just be the solutions to $\hat{H}_0$, which we assume we know, and $\lambda = 1$ represents the solution to the full problem. However, in principle, we can also solve the problem for any intermediate λ, and so we can assume that any quantity $O(\lambda)$ we are interested in will be a continuous function of λ. We can then expand the solution in a power series of λ

$$O(\lambda) = O^{(0)} + O^{(1)}\lambda + O^{(2)}\lambda^2 + \cdots = \sum_n O^{(n)}\lambda^n \tag{2.17}$$

The trick is just to have an efficient scheme to calculate the series coefficients $O^{(n)}$.

One thing worth mentioning, however, is that in QED the perturbative series of the form (2.17) is an asymptotic expansion and the series does not formally converge [53]. While this often does not matter in practice as the asymptotic expansion generally replicates experiments extremely well up to high order in the expansion, it does mean that perturbation theory for relativistic QFTs is fundamentally incomplete.

Time-dependent perturbation theory and the S-matrix

In the case of time-dependent perturbation theory, we are interested in calculating the probability of observing the system in a final state $|f\rangle$ at time t, given that the system started in an initial state $|i\rangle$ at time t_0. In quantum mechanics, we can connect initial and final states of a system with the unitary time-evolution operator $U(t, t_0)$, such that the state of the system $|\Psi(t_0)\rangle$ at time t_0 propagating forward in time is

$$|\Psi(t)\rangle = U(t, t_0)\,|\Psi(t_0)\rangle$$

[21] Or the Lagrangian. It goes beyond the scope of this text, but it is also possible to formulate QFT and perturbation theory in the *path integral formalism*, where the core quantities under consideration are the Lagrangian and the action. [8]

Assuming that the time-evolution is generated by a general time-dependent Hamiltonian $H(t)$, it is possible to write down the time-evolution operator as the infinite series

$$\begin{aligned} U(t,t_0) &= 1 - i\int_{t_0}^{t} H(t_1)\,\mathrm{d}t_1 + (-i)^2 \int_{t_0}^{t}\int_{t_0}^{t_1} H(t_1)H(t_2)\,\mathrm{d}t_2\,\mathrm{d}t_1 + \cdots \\ &= \sum_{n=0}^{\infty}(-i)^n \int_{t_0}^{t}\int_{t_0}^{t_1}\cdots\int_{t_0}^{t_{n-1}} H(t_1)\cdots H(t_n)\,\mathrm{d}t_n\cdots\mathrm{d}t_1 \end{aligned} \tag{2.18}$$

which is commonly referred to as the *Dyson series* [54]. In the special case of a time-independent Hamiltonian, the Dyson series is equal to the well-known exponential form of the time-evolution operator

$$U(t,t_0) = \exp\left[-iH(t-t_0)\right]$$

It is formally possible to recover the exponential form in the case of a general time-dependent Hamiltonian as well, and the general time-evolution operator can be written as

$$U(t,t_0) = T\exp\left(-i\int_{t_0}^{t} H(\tau)\,\mathrm{d}\tau\right) = \sum_{n=0}^{\infty}\frac{(-i)^n}{n!}\int_{t_0}^{t}\cdots\int_{t_0}^{t} T\left[H(t_1)\cdots H(t_n)\right]\mathrm{d}t_n\cdots\mathrm{d}t_1$$

where the time-ordering meta-operator T ensures that the time coordinates in the product of the Hamiltonians are always in decreasing order, permuting them around as necessary. The second step, where the exponential has been expanded in its Taylor series and the time-ordering moved into the integrand, is formally the definition of the time-ordering of an exponential operator. Crucially, compared to (2.18), there is no more dependence on the integration variables in the integration bounds. This form is useful because a time-independent Hamiltonian (i.e. physically time-independent, or in other words, lacking explicit time dependence in the Schrödinger picture) acquires a time-dependence in the Heisenberg and interaction pictures of quantum mechanics, and perturbation theory is usually formulated in terms of the latter.

Finally, in perturbative problems, the initial and final times are usually formally allowed to approach infinity. This leads to the *scattering matrix* S_{fi}, which is simply the transition amplitude from the initial state to the final state over an infinite time interval

$$S_{fi} = \lim_{\substack{t\to+\infty \\ t_0\to-\infty}} \langle f|U(t,t_0)|i\rangle = \langle f|\ T\exp\left(-i\int_{-\infty}^{\infty} H(\tau)\,\mathrm{d}\tau\right)|i\rangle \tag{2.19}$$

and is the quantity that is commonly used in perturbative calculations.

Feynman diagrams

When the Dyson series (2.18) is inserted into the expression (2.19) for the scattering matrix, it becomes clear that the scattering matrix can be evaluated as an infinite series, which constitutes the perturbative expansion.

In order to separate out the time-evolution of the non-interacting "easy" part, and the "hard" interaction parts, the calculation is generally done in the so-called *interaction picture* of quantum mechanics, where the time-evolution generated by the free part of the Hamiltonian moved from states to operators

$$|\Psi\rangle_{\mathrm{I}} = \mathrm{e}^{iH_0 t}\,|\Psi\rangle\,, \quad A_{\mathrm{I}}(t) = \mathrm{e}^{iH_0 t}A(t)\mathrm{e}^{-iH_0 t}$$

where $|\Psi\rangle_{\mathrm{I}}$ and $A_{\mathrm{I}}(t)$ are the states and operators in the interaction picture. The practical consequence is that the Hamiltonian in the Dyson series expansion reduces to just the interaction part $\hat{H}_{\mathrm{I}}$ of the Hamiltonian.

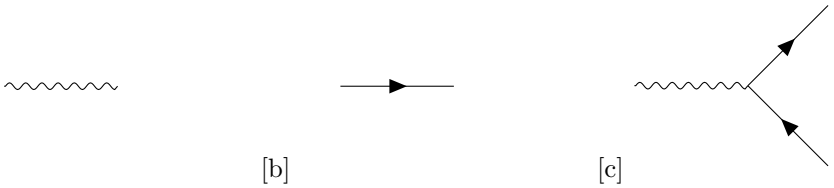

[a] [b] [c]

Figure 2.2: The building blocks of QED Feynman diagrams: (a) the photon propagator, (b) the electron-positron propagator and (c) the QED interaction vertex.

In the interaction picture, the integrals for each of the terms in the perturbative expansion can be constructed with the help of Feynman diagrams, which are a type of graph. Each vertex of such a graph corresponds to an interaction term, with the number of vertices giving the order of the diagram. Generally, each interaction term brings some coupling constant which, if small, can be used to argue for the convergence of the series expansion — higher-order terms have a higher-order coupling constant, hopefully leading the series to converge. The edges of the Feynman diagrams are called *propagators* and are related to the time-evolution of the free particle states under the free part of the Hamiltonian. The propagators can be thought of as either Green's function of the free part of the Hamiltonian, or as certain vacuum expectation values of the field operators [24].

The three basic building blocks of free particle QED diagrams are shown in figure 2.2. In QED there is only one interaction vertex with a coupling coefficient α, the fine-structure constant. The diagrams are then classified according to the number of vertices or, equivalently, in powers of α, with higher-order diagrams ($\sim \alpha^n$) usually contributing less to whichever process one happens to study.

Energy shifts of stationary states

This thesis focuses on stationary bound states, rather than time-dependent or scattering processes. As such, it would be more natural to employ time-independent perturbation theory, which would lead to a series expansion similar to Eq. (2.17), where the terms are calculated via matrix elements of the interaction operator in the basis of free-particle solutions [55]. However, it turns out that the methods of time-dependent perturbation theory are very much applicable in the relativistic case too.

It was shown by Gell-Mann and Low in 1951 (the Gell-Mann-Low theorem) [56, 57] that the formal perturbative solution for the ground state of an interacting system $|\Psi_0\rangle$

$$\mathcal{N}\,|\Psi_0\rangle = [1 + (H_0 - \epsilon_0)^{-1}(1 - |0\rangle\,\langle 0|)(H_\mathrm{I}(0) - E_0 + \epsilon_0)]^{-1}\,|\Phi\rangle$$

given relative to the vacuum state $|0\rangle$ acting as the ground state of the free part of the Hamiltonian H_0, can be written down in terms of the time-evolution operator $U(t, t_0)$ such that

$$\mathcal{N}\,|\Psi_0\rangle = \frac{U^{-1}(\pm\infty, 0)\,|0\rangle}{\langle 0|U^{-1}(\pm\infty, 0)|0\rangle}$$

Here $\mathcal{N}$ is a normalization constant that would have to be determined separately. This, in turn, leads to the expression for the energy shift of an arbitrary eigenstate $|n\rangle$ of the free-particle Hamiltonian

$$\Delta E_n = \lim_{\substack{\varepsilon\to 0\\ \lambda\to 1}} \frac{\langle n|U(\infty,0)[H_0 + \lambda H_\mathrm{I}^\varepsilon(0) - E_n]U(0,-\infty)|n\rangle}{\langle n|U(\infty,-\infty)|n\rangle} = \lim_{\substack{\varepsilon\to 0\\ \lambda\to 1}} \frac{i\varepsilon\lambda}{2}\frac{\frac{\partial}{\partial\lambda}S^\varepsilon_{nn}}{S^\varepsilon_{nn}}$$

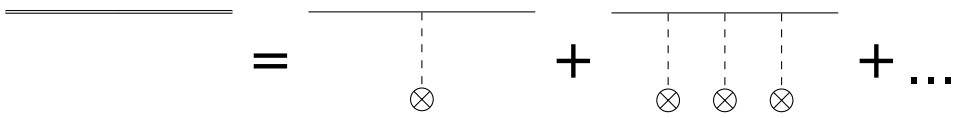

Figure 2.3: The bound state electron propagator is conventionally represented by a double line. It can be further be split up as a sum in orders of $Z\alpha$, where each "interaction" with the nuclear potential is marked by a dashed line and a cross. It should be emphasized here that the dashed line does not represent a propagator and that the expansion in terms of $Z\alpha$ only works if $Z\alpha \ll 1$.

which can also be written down in terms of the S-matrix. The additional ε parameter is from the formal use of an adiabatically damped interaction Hamiltonian $H_{\mathrm{I}}^{\varepsilon}(t) = \mathrm{e}^{-\varepsilon|t|} H_{\mathrm{I}}(t)$, which allows certain infinities in the perturbative expansion to be correctly cancelled out. [58, 59]

Hence, the energy shifts due to the QED interaction, which is a time-independent problem, can be evaluated with the help of the scattering matrix. As a consequence, it is possible to apply the methods of time-dependent perturbation theory, such as Feynman diagrams, to the problem of bound state energy shifts.

Furry picture for bound states

While it is definitely possible to use non-interacting plane wave states as the reference states in perturbative bound-state calculations (i.e. using the free-particle Hamiltonian as H_0) [60], it can be advantageous, for convergence reasons, to include the electromagnetic scalar potential $V(\boldsymbol{x})$, such as the one generated by the nucleus, in the unperturbed part as well, i.e.

$$H_0 = h_{\mathrm{D}} + V(\boldsymbol{x})$$

where h_{D} would be the free-particle Dirac Hamiltonian. This avoids the need to solve the bound-state problem perturbatively, which is good since the bound states are qualitatively very different from the plane-wave states. This leads to the *Furry picture* of perturbative bound state QED. [61]

In practice, this means that the reference "free" states are in fact the bound states solutions of the external field Dirac equation. The effect of the potential on the time-evolution is incorporated into the propagator, and such propagators are commonly depicted with double lines. As illustrated in Figure 2.3, the bound state propagators can be further expanded such that they become a product of the potential and the free electron propagators in various orders [49, Eq (14-70) and (14-71)]. In the case of a potential generated by a nucleus[22], each term in this expansion has a factor of $(Z\alpha)^n$, and so the Feynman diagrams can be classified first in orders of α (due to the electromagnetic interaction terms) and then in orders of $(Z\alpha)$ (due to the nuclear potential). Often simply considering the diagrams in lowest order of $(Z\alpha)$, is sufficient, but it does break down in the regime where $Z\alpha \sim 1$, which is the case e.g. for the super-heavy elements.

Effective Hamiltonians

A way to avoid having to treat a fully quantized photon field is to replace it with some effective operators that approximate the interaction as accurately as possible.

[22]The fine-structure constant α is related to the elementary charge by $\alpha \propto e^2$. Hence, the nuclear potential of a nucleus with charge $Q = Ze$ will be proportional to $V \propto eQ = Ze^2 \propto Z\alpha$.

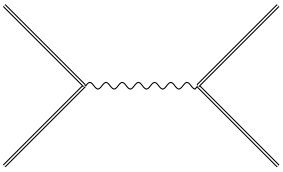

Figure 2.4: Feynman diagram of the one-photon exchange that leads to the Breit interaction correction of the electron-electron interaction.

Electron-electron interactions

Most significant of such replacements is the introduction of a two-electron interaction, to describe the effect one electron has on another via the photon field (or a positron on a positron, or an electron on a positron). It should be pointed out that the full QED Hamiltonian does not contain any terms for direct interaction between electrons, and the interaction is mediated by the photon component of the wavefunction instead. The closest one gets is the electromagnetic interaction term which contains three (as opposed to four, of a standard two-particle interaction term) creation-annihilation operators and couples the photon field to individual electrons and positrons, or the creation of annihilation of their pairs.

A formal way to derive such an effective operator from QED is to derive the kernel $g(R;\omega)$ of the matrix element of the one-photon exchange from perturbation theory, with the diagram illustrated in Figure 2.4. [62, Section 4.9] In the Coulomb gauge, this leads to the kernel

$$g(R;\omega) = \frac{1}{R} + g^{\mathrm{T}}(R;\omega)$$

where non-Coulomb part is given by the *transverse photon kernel*

$$g^{\mathrm{T}}(R;\omega) = -\boldsymbol{\alpha}_1 \cdot \boldsymbol{\alpha}_2 \frac{e^{i\omega R/c}}{R} - (\boldsymbol{\alpha}_1 \cdot \nabla)(\boldsymbol{\alpha}_2 \cdot \nabla)\frac{e^{i\omega R/c} - 1}{\omega^2 R} \tag{2.20}$$

As can be seen, if the transverse photon part is removed, it reduces to the well-known instantaneous Coulomb interaction kernel. The operator (2.20) is sometimes referred to as the *frequency-dependent Breit interaction*. In the long wavelength limit, where $\omega \to 0$, it reduces down to the more well-known *Breit interaction*

$$g^{\mathrm{B}}(R;\omega) = \lim_{\omega\to 0} g^{\mathrm{T}}(R;\omega) = \frac{1}{2R}\left(\boldsymbol{\alpha}_1 \cdot \boldsymbol{\alpha}_2 + (\boldsymbol{\alpha}_1 \cdot \hat{\boldsymbol{R}})(\boldsymbol{\alpha}_2 \cdot \hat{\boldsymbol{R}})\right) \tag{2.21}$$

where $\hat{\boldsymbol{R}} = \boldsymbol{R}/R$ is a unit vector. The "frequency-independent Breit" interaction was originally derived using a more classical, non-QED approach by considering two moving charges within a relativistic framework [63].

As a final point, it should be pointed out that depending on the choice of gauge, the one-photon exchange diagram in Figure 2.4 can lead to different effective Hamiltonian. Specifically, in the Feynman gauge, the matrix element kernel becomes

$$g^{\mathrm{M};\omega} = (1 - \boldsymbol{\alpha}_1 \cdot \boldsymbol{\alpha}_2)\frac{\mathrm{e}^{i\omega R/c}}{R}$$

which is known as the *Møller interaction*. In the long-wavelength limit $\omega \to 0$ it reduces to

$$\lim_{\omega\to 0} g^{\mathrm{M}} = \frac{1 - \boldsymbol{\alpha}_1 \cdot \boldsymbol{\alpha}_2}{R}$$

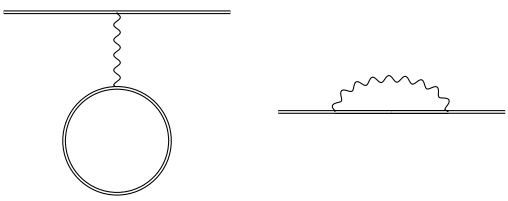

Figure 2.5: Furry picture Feynman diagrams for vacuum polarization and self-energy in first-order in α.

where the first term is again the non-relativistic Coulomb interaction and the second term

$$g^{\mathrm{G}} = \frac{-\boldsymbol{\alpha}_1 \cdot \boldsymbol{\alpha}_2}{R}$$

is known as the *Gaunt interaction* [64], which contains spin-other-orbit interactions and is commonly used as a correction to the Coulomb interaction in relativistic electronic structure calculations of molecules.

Vacuum polarization

Another correction QED gives rise to has to do with the creation of electron-positron pairs by the electromagnetic field generated by the nucleus. This leads to an induced charge distribution near the nucleus or, in other words, the nucleus *polarizes the vacuum*. This is directly related to the pair creation and annihilation in the term that couples the external electromagnetic field to the quantized Dirac field, as was discussed in Section 2.4.

The charge distribution resulting from vacuum polarization can be calculated in perturbation theory from the vacuum expectation value of the current operator

$$j^\mu = e \langle 0 | \bar{\psi} \gamma^\mu \psi | 0 \rangle$$

Only the time-like component (i.e. the charge distribution) ρ_{vp} is non-zero. The induced charge will then, in turn, induce an electric potential which will modify the Coulomb potential of the nucleus. If the charge distribution is known, the corresponding potential can, with relative ease, be determined from the Poisson equation

$$\nabla^2 \varphi_{\mathrm{vp}}(\boldsymbol{x}) = \rho_{\mathrm{vp}}(\boldsymbol{x})$$

This potential can then be incorporated into electronic structure calculations simply as a modification of the nuclear potential.

The Feynman diagram corresponding to the lowest order vacuum polarization in the Furry picture can be seen in Figure 2.5. Both in practice and historically, that contribution is discussed in two parts, which can be linked to the $(Z\alpha)$ expansion of the diagram, illustrated in Figure 2.6. The first term of the $(Z\alpha)$ expansion leads to the *Uehling contribution*, originally considered by Uehling in the 1930s, way before QED was invented [65]. The other terms, in orders $(Z\alpha)^{3+}$, were calculated much later in the 1950s by Wichmann and Kroll [66].

The vacuum polarization also has higher-order contributions. Källén and Sabry considered the contribution of vacuum polarization in second order[23] in α [60]. In terms of the Furry

[23]The relevant paper is called "Fourth Order Vacuum Polarization", which refers to the fact that second order in α corresponds to fourth order in the elementary charge e, as $\alpha \propto e^2$.

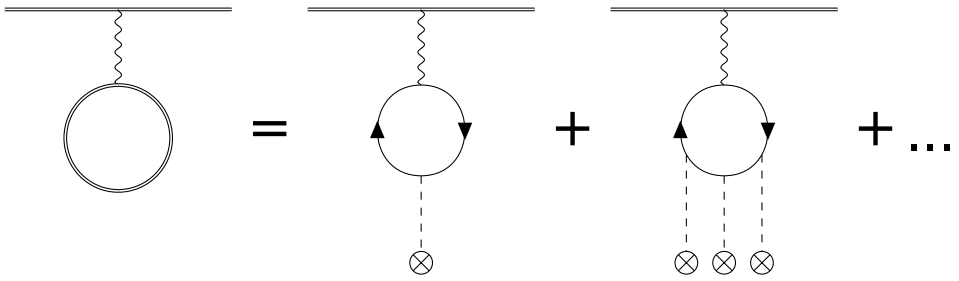

Figure 2.6: Expansion of the VP diagrams in orders of $Z\alpha$. Note that the dashed lines are not really propagators but part of a single vertex, each corresponding to an order $Z\alpha$. Note that Furry's theorem only gives odd couplings to the external potential (dashed lines).

picture, the expressions derived by Källén and Sabry are only in first order in $Z\alpha$ [59], which is due to the zeroth-order Hamiltonian being the free-particle Hamiltonian in Källén's formulation of perturbative QED [67]. In that sense it is similar to the Uehling term in the first-order vacuum polarization.

Self-energy

Another important effect introduced by QED is that an electron can, essentially, interact with itself. Physically one can think of this as an electron emitting a photon at a point in time and then reabsorbing it later, which leads to an energy shift called the *self-energy*.

Historically, the self-energy was important in explaining the Lamb shift — the experimental energy splitting in single-electron states with the same principal quantum number n and total angular momentum j, observed by Lamb and Retherford in the 1940s. [43] According to the Dirac equation these states should be energetically degenerate, and the vaccuum polarization contribution, while significant at higher nuclear charge, is too small and has the wrong sign, and hence can not explain the shift either. Self-energy as an explanation for this shift was first considered by Bethe in 1947 [44], although for non-relativistic electrons, with relativistic arguments only used to argue for a cutoff value in the energy of the photon propagator. In 1949, this was generalized to full relativistic theory by Kroll and Lamb [45], but in either case it showed very good agreement with the experiment, indicating that the Lamb shift can be explained by considering the electron's self-interaction interaction via the radiation field.

In the Furry picture, the self-energy correction can be considered in the lowest order of the electromagnetic coupling α via the diagram shown in Figure 2.5. A way to calculate the self-energy integral without performing a series expansion in $Z\alpha$ was shown by Mohr in 1974 [68], effectively summing the $Z\alpha$ expansion analytically across all terms. He then went on to evaluate this integral numerically for the $1s_{1/2}$ state of the neutral $Z = 10, 20, \ldots, 110$ atoms in the follow-up article [69]. Furthermore, in [70] he discusses the generalization of his approach to all principal quantum number $n = 2$ states and published a table of numerically evaluated energy shifts for $Z = 10, \ldots, 40$ for all the $n = 1, 2$ states in [71]. Finally, in [72], the self-energy calculations were once again extended to all $n \leq 5$, $j = \frac{1}{2}, \frac{3}{2}$ states, and the shifts were evaluated for $Z = 10, \ldots, 110$.

Mohr's approach only determines the energy shifts and just for hydrogenic-like systems. By replacing the nuclear charge Z by some screened effective charge Z_{eff} in multi-electron systems, these can be used to perturbatively estimate the self-energy correction for a many-electron system. However, an operator form, similar to the Breit interaction or vacuum polarization

potential, would be preferred. As such, over the years, various effective operators [18, 19, 20] and other approaches [73, 74] have been proposed. Incorporating self-energy corrections into many-electron calculations of large systems is tricky, and the focus of the work in Chapter 4. That also includes a review of the aforementioned effective operators.

Higher-order corrections

The effects described in the preceding sections by no means comprehensively describe the effect of the fully quantized photon field. It would also be possible to consider higher-order vacuum polarization and self-energy terms. However, the corrections arising from including these additional diagrams are generally assumed to be small and only required in extremely high-precision calculations of few-electron systems. They are briefly reviewed in the 1998 review by Mohr, Plunien and Soff [59].

2.6 Spherical symmetry and radial external-field Dirac equation

Atoms, consisting of electrons captured by a single atomic nucleus, are to a very good approximation spherically symmetric systems.[24] The charge distributions of the atomic nuclei (see also Section 2.7) do not experience any forces that would break spherical symmetry, and consequently neither do the electron clouds. This can be taken advantage of to considerably reduce the number of variables and degrees of freedom when analysing these systems.

In practice, this generally means that after reformulating the problem in spherical coordinates, the angular part can be separated out and solved analytically. The non-trivial radial part will become a one-dimensional problem on the $r \in [0, \infty)$ axis. It also leads to conserved quantities and quantum numbers in the form of angular momenta that are useful in labelling atomic states. This section briefly discusses general angular momentum theory and how this is applied when reducing the single-particle Dirac equation to reduce to a radial equation, while the application of spherical symmetry to the many-electron problem is discussed in Section 4.1.

Angular momentum theory

Spherical symmetry means that the system (i.e. the Hamiltonian) is invariant under rotation in three dimensional space around a point. This is true for an atom as long as the nucleus is centered at the origin of the coordinate system. An extremely useful way of describing spherical symmetry and angular momentum in quantum mechanics is to connect the description to the group theory of rotation in three dimensions, with the group elements corresponding to angular momentum operators.[25]

In the abstract group theoretical language, the operators $R(\hat{\boldsymbol{n}}, \phi)$, effecting a rotation of angle ϕ around an axis described by the unit vector $\hat{\boldsymbol{n}}$, form a (continuous) group.[26] It is possible to

[24] While the electromagnetic interaction in QED does not break spherical symmetry, the parity-violating weak interaction does. This means that even if all the external forces are perfectly spherical, atomic systems fundamentally do not obey spherical symmetry fully, and even an isolated electron has, in principle, a tiny electric dipole moment. However in practice, the parity violation effects of the weak interaction are extremely small and can be neglected or treated perturbatively.

[25] The books by Edmonds and Rose are excellent brief, concise overviews of angular momentum theory. [75, 76] If something more substantial is needed, the book by Varshalovich, Moskalev & Khersonskiĭ is an excellent reference. [77]

[26] In the more general mathematical abstraction such continuous transformation form a *Lie group* and the generalization of the commutation relations of the generators is known as a *Lie algebra*. In particular, spin and rotational angular momentum are related to the representation of the two-dimensional *special unitary group* $\mathrm{SU}(2) \subset \mathrm{GL}(2, \mathbb{C})$ which in turn contains the three-dimensional real *special orthogonal group* $\mathrm{SO}(3) \subset \mathrm{GL}(3, \mathbb{R})$

define the corresponding *infinitesimal generator* of rotations around the $\hat{\boldsymbol{n}}$-axis as

$$J(\hat{\boldsymbol{n}}) = i\hbar \frac{\partial R(\hat{\boldsymbol{n}}, \phi)}{\partial \phi}\bigg|_{\phi=0} \tag{2.22}$$

where $\hbar$ is introduced to give the quantity the appropriate physical units. It makes sense to think of the set of generators $J(\hat{\boldsymbol{n}})$ as living on an unit sphere in the three dimensional space, and so, while somewhat of an abuse of notation, it is convenient to think of this set $\boldsymbol{J}$ being represented by a single three-vector operator, which can be projected down to a particular axis $\hat{\boldsymbol{n}}$ via the scalar product

$$J(\hat{\boldsymbol{n}}) = \hat{\boldsymbol{n}} \cdot \boldsymbol{J}$$

By effectively inverting (2.22), the rotation operator can be written down in an exponential notation in terms of the generators as

$$R(\hat{\boldsymbol{n}}, \phi) = \exp\left(-\frac{i\phi}{\hbar}\hat{\boldsymbol{n}} \cdot \boldsymbol{J}\right) \tag{2.23}$$

Commonly, there is an implicit or explicit orthogonal coordinate system and it is convenient to just consider the generators along the x, y and z axes

$$\boldsymbol{J} = (J_x, J_y, J_z)$$

which form a basis for representing the general generators as

$$J(\hat{\boldsymbol{n}}) = \hat{\boldsymbol{n}} \cdot \boldsymbol{J} = \sum_i \hat{n}_i J_i$$

These angular momentum operators in an orthogonal coordinate system obey commutation relations with each other

$$[J_i, J_k] = i\hbar \sum_k \varepsilon_{ijk} J_k \tag{2.24}$$

when written in terms of the three-dimensional Levi-Civita symbol ε_{ijk}.

A key insight from the commutator in Eq. (2.24) is that the three operators along the orthogonal axes do not form a commuting set and therefore can not be used to label states by their eigenvalues even if the system has a well-defined angular momentum. Instead it is useful to consider the square of the total angular momentum[27]

$$\boldsymbol{J}^2 = J_x^2 + J_y^2 + J_z^2$$

which commutes with every angular momentum operator

$$[\boldsymbol{J}^2, J_i] = 0$$

Its eigenvalue together with the eigenvalue of one of the angular momentum operators, conventionally J_z, can then be used to label states. The eigenvalues are defined by

$$\boldsymbol{J}^2 |\Psi\rangle = j(j+1) |\Psi\rangle , \quad J_z |\Psi\rangle = m |\Psi\rangle$$

describing spacial rotations. For a more general introduction to how the mathematical theory of groups is applied in science, Ludwig & Falter is a very accessible resource. [78]

[27] In the Lie group language, this is the unique *the Casimir invariant* of the SO(3) group.

and it can be shown, purely from the algebra, that they only take certain discrete values. The total angular momentum of a state can only adopt positive integer or half-integer values

$$j = 0, 1/2, 1, 3/2, \dots$$

and for a given total angular momentum j, the m quantum number can only take the $2j + 1$ distinct values

$$m = j, j-1, j-2, \dots, -j+1, -j$$

changing in integer steps between j and $-j$. For a half-integer total angular momentum, $m = 0$ is not a valid quantum number.

Rotational invariance

Mathematically, spherical symmetry can be defined as rotational invariance, i.e. that an object or a system does not change under a rotation. For an operator A, such as the Hamiltonian of a system, and with the help of the rotation operator (2.23), this condition can be written as

$$\forall \boldsymbol{\phi} : \exp(-i\boldsymbol{\phi} \cdot \boldsymbol{J}) A \exp(i\boldsymbol{\phi} \cdot \boldsymbol{J}) = A$$

where $\boldsymbol{\phi} = \phi \hat{\boldsymbol{n}}$ is a rotation of angle ϕ around $\hat{\boldsymbol{n}}$. This is equivalent to requiring the operator to commute with the angular momentum operators, and consequently with the total angular momentum operator $\boldsymbol{J}^2$

$$[A, J_k] = 0, \quad [A, \boldsymbol{J}^2] = 0$$

The eigenstates of the system can then be labelled using the angular momentum quantum numbers. Another way to think about it is that the operator can be block-diagonalized into angular momentum subspaces.

Physical sources of angular momentum

In a physical system, there can be many separate sources of angular momentum.

- The most significant is simply the rotational angular momentum $\boldsymbol{L}$ of a particle, which corresponds to the actual physical rotation of the system or particle around the origin of the coordinate system. In a multi-particle system, every particle brings its own rotational angular momentum and corresponding operators, which all add up to a total angular momentum. However, it should be noted that in a quantum field theory picture, particles of the same species can be thought of as excitations of a single underlying field, and as such we would be talking about the angular momentum of a field. However, there would still be a angular momentum operator for each separate field.
- Elementary particles also carry inherent, internal angular momenta: their spin $\boldsymbol{S}$. As such, fermionic particles such as electrons and quarks, have an additional internal spin angular momentum of 1/2. Similarly, vector bosons such as photons or the Z- or W-bosons, are spin-1 particles, and the hypothetical gravitons would be spin-2. The Higgs boson, on the other hand, is a massive spin-0 scalar particle and does not carry any inherent spin.
- In atomic physics it is also common to talk about the total spin of the atomic nucleus: the nuclear spin $\boldsymbol{I}$, which we generally assume to be a known, given quantity. However, it should be pointed out that, in principle, the nuclear spin itself is made up of the rotational and internal angular momenta of the constituent particles (the quarks and the gluons).[28]

[28] The question of how spin arises in protons and neutrons is still hotly debated [79].

All the different sources of angular momenta add up. For atoms, in most cases, we are only concerned with the total angular momenta of the electrons

$$\boldsymbol{J} = \boldsymbol{L} + \boldsymbol{S} \tag{2.25}$$

made up of their rotational and spin angular momenta. In the case where we are also interested in the effect the nucleus has on the electronic structure, we add the nuclear spin

$$\boldsymbol{F} = \boldsymbol{J} + \boldsymbol{I} = \boldsymbol{L} + \boldsymbol{S} + \boldsymbol{I}$$

to get the total angular momentum of the whole atom $\boldsymbol{F}$.

Coupling of angular momenta

It would be tempting to think that, if we combine the angular momenta from two sources, say $\boldsymbol{L}$ and $\boldsymbol{S}$, and the state is an eigenstate of $\boldsymbol{L}^2$, L_z, $\boldsymbol{S}^2$ and S_z, that it would then also be an eigenstate $|\ell, m_\ell, s, m_s\rangle$ of the corresponding operators $\boldsymbol{J}^2$ and J_z of the total angular momentum $\boldsymbol{J} = \boldsymbol{L} + \boldsymbol{S}$. However, in general, that is not the case. While the m quantum number does add to $m_j = m_\ell + m_s$, this does not hold for the total angular momentum quantum number.

Instead, it is necessary to consider the whole subspace spanned by all the m_ℓ and m_s quantum numbers for given ℓ and s. Then, via certain superpositions, it is possible to construct states that are also eigenstates of J_z

$$|\ell, s; j, m_j\rangle = \sum_{m_\ell, m_s} C(j, \ell, s; m_j, m_\ell, m_s)\, |\ell, m_\ell, s, m_s\rangle \tag{2.26}$$

with valid j values limited to $j = |\ell + s|, |\ell + s| - 1, \ldots, |\ell - s| + 1, |\ell - s|$. The coefficients $C(J, j_1, j_2; M, m_1, m_2)$ are known as the Clebsch-Gordan coefficients and are the same for all physical systems, as they can be derived directly from the abstract algebraic properties of the angular momentum operators. The coefficients also exhibit various symmetries, such as the requirement that $M = m_1 + m_2$, which can be interpreted as a symmetry where the coefficient is zero if that condition is not met.

The states resulting from Eq. (2.26) are now eigenstates of all three angular momentum operators $\boldsymbol{J}^2$, $\boldsymbol{L}^2$ and $\boldsymbol{S}^2$. However, while an eigenstate of J_z, those states are no longer eigenstates of L_z nor S_z. An important point here is that this coupling can be thought of as an basis transformation within each ℓ and s subspace.

κ quantum number

When coupling the orbital and spin angular momenta into total angular momentum in line with Eq. (2.25), the total angular momentum quantum number j is no longer sufficient to uniquely identify the coupled states and we have to keep track of both ℓ and j quantum numbers. This is because for each integer orbital angular momenta we get two states with $j = \ell \pm 1/2$, which means that for each valid half-integer j value, we always have two states with the same total angular momentum, e.g. $s_{1/2}$ and $p_{1/2}$, or or $p_{3/2}$ and $d_{3/2}$.

As such, in relativistic theory it is convenient to use the κ quantum number, which can encode the orbital and total angular momentum labels as a single integer value. The relationship between the three quantum numbers can be summarised as

$$\kappa = \mp\left(j + \frac{1}{2}\right), \quad \text{for} \quad j = \ell \pm \frac{1}{2}$$

and it essentially just assigns either a positive or a negative integer for distinguishing the two orthogonal functions with the same total angular momentum. This is similar to how the ℓ quantum number can distinguish between two different coupled states in non-relativistic theory where the total angular momentum eigenstates are constructed by explicitly coupling orbital and spin angular momentum eigenstates. In fact, the ℓ-value based nomenclature is still used in relativistic theory, but the difference here is that the Dirac states are not actually eigenstates of the orbital angular momentum.

The κ quantum number can also be defined as the eigenvalue of the "spin-orbit operator"

$$K = -\beta(2\boldsymbol{S}\cdot\boldsymbol{L}+1) = \beta\left(\boldsymbol{L}^2 - \boldsymbol{J}^2 - \frac{1}{4}\right)$$

Some angular momentum eigenstates $|n\kappa m\rangle$ would therefore have eigenvalues

$$\boldsymbol{J}^2|n\kappa m\rangle = j(j+1)|n\kappa m\rangle\,, \quad J_z|n\kappa m\rangle = m|n\kappa m\rangle\,, \quad K|n\kappa m\rangle = \kappa|n\kappa m\rangle$$

The operators $\hat{\boldsymbol{J}}$ and $\hat{K}$ commute with each other, and with the Dirac-Couloumb Hamiltonian.[29]

Radial Dirac equation

A special form of the Dirac equation, corresponding to the external field Dirac Hamiltonian (2.14) with a scalar time-like time-independent radially symmetric central potential $V(r)$ such that $A^\mu(x) = (V(r), \mathbf{0})$ is

$$\left[i\gamma^\mu\partial_\mu - m + e\gamma^0 V(r)\right]\psi(x) = 0 \tag{2.27}$$

With a spherically symmetric external potential, the single-particle eigenstates will simultaneously be total angular momentum eigenstates too. Following the convention[30] in Grant's book [62], the radial orbitals can be written as

$$\psi(\boldsymbol{x}) = \frac{1}{r}\begin{pmatrix} P(r)\chi_{\kappa m}(\theta,\varphi) \\ iQ(r)\chi_{-\kappa m}(\theta,\varphi)\end{pmatrix} \tag{2.28}$$

where $P(r)$ and $Q(r)$ are single variable complex functions on $r \in [0,\infty)$ and $\chi_{\kappa m}(\theta,\varphi)$ are the spin-spherical harmonics (see Appendix A). The Dirac equation then transforms into a set of two coupled first-order differential equations, which, in full SI units, can be written as a two-component matrix differential equation

$$\begin{pmatrix} mc^2 + eV(r) & \hbar c\left(-\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right) \\ \hbar c\left(\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right) & -mc^2 + eV(r)\end{pmatrix}\begin{pmatrix} P(r) \\ Q(r)\end{pmatrix} = E\begin{pmatrix} P(r) \\ Q(r)\end{pmatrix} \tag{2.29}$$

Especially when implementing a numerical solver, it is very useful to further simplify the equation by choosing an appropriate set of natural units. In Hartree atomic units ($\hbar = e = a_0 = m_e = 1$, $c = \alpha^{-1}$) the equation becomes

$$\begin{pmatrix} m\alpha^{-2} + V(r) & \alpha^{-1}\left(-\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right) \\ \alpha^{-1}\left(\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right) & -m\alpha^{-2} + V(r)\end{pmatrix}\begin{pmatrix} P(r) \\ Q(r)\end{pmatrix} = E\begin{pmatrix} P(r) \\ Q(r)\end{pmatrix} \tag{2.30}$$

[29] A more detailed exposé of the exact definition and relationship of rotation and the Dirac spinors can be found in Chapter 4.6 of Thaller's book on the Dirac equation [11].

[30] Different authors use different sign conventions, especially for the $Q(r)$ function, so care must be taken when comparing expressions from different sources.

where the mass is still kept as a separate parameter, but will be given in units of electron mass.

However, from both a mathematical and numerical point of view, a more convenient system of units are the particle physics natural units ($\hbar = m_e = c = \epsilon_0 = 1$) where the fine-structure constant is directly related to the elementary charge via $e = \sqrt{4\pi\alpha}$. If the nuclear potential is written in terms of a "normalized" potential $\Phi(r)$ as

$$V(r) = \frac{eZ}{4\pi\epsilon_0}\Phi(r) \tag{2.31}$$

then the Dirac equation can be written as

$$\begin{pmatrix} m + (Z\alpha)\Phi(r) & -\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r} \\ \frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r} & -m + (Z\alpha)\Phi(r) \end{pmatrix} \begin{pmatrix} P(r) \\ Q(r) \end{pmatrix} = E \begin{pmatrix} P(r) \\ Q(r) \end{pmatrix} \tag{2.32}$$

In this form it becomes very clear that $Z\alpha$ can be seen as a coupling constant between a normalized (corresponding to a unit charge) nuclear charge density and the electrons.

Atomic units: an uncertain choice

There is a very pragmatic reason to choose the particle physics natural units over the atomic units when working with the Dirac equation. It atomic units, the scale of the energy will be proportional to α^{-2}, where α is the fine-structure constant, with the value [23]

$$\alpha = 7.2973525693(11) \times 10^{-3}$$

What is important here is that the value of the fine-structure constant is experimentally derived and therefore has an experimental uncertainty attached to it, which currently is on the order of $u(\alpha) \sim 10^{-12}$. As the energy scale in atomic units is proportional to α^{-2} (i.e. $E \sim \alpha^{-2}$ Ha), this leads to an inherent uncertainty of

$$u(E) \sim 2\alpha^{-3}u(\alpha) \sim 10^{-5} \text{ Ha}$$

of the energy. Any numerical result reported today that is reported in atomic units will inherently be limited by this uncertainty. It should also be noted that as the energy scale is 10^4 Ha and the machine precision of double precision is 10^{-16}, we would expect that we can achieve a numerical precision of 10^{-12} Ha, much lower than the one imposed by the fine-structure constant.

It should also be noted that this issue not theoretical, but actually arose when working on Chapter 3. At one stage a slight inconsistency of $\delta\alpha = 2.9 \times 10^{-12}$ in the value of α — one part of the code was using the CODATA 2014 [80], whereas the other part of the code had been updated to use the CODATA 2018 [23] value — lead to a systematic discrepancy in the energies of exactly $\delta E = 2\alpha^{-3} \cdot \delta\alpha = 1.49 \times 10^{-5}$. This error would have been avoided if the code and the intermediate values would have been using natural units.

Mass scaling

It is possible to simplify the Equation (2.32) even further by eliminating the mass from the equation. For that, the radial coordinate needs to be rescaled via $x = rm$ and the normalized potential must have a specific scaling behaviour

$$\Phi(x/m) = m\Phi(x)$$

This clearly holds for the Coulomb potential $\Phi(r) = 1/r$, but will actually hold for any reasonable charge distribution. This can be argued for through dimensional analysis: the dimension of Φ

is L^{-1} (inverse length). As long as all the parameters of the potential that have the dimension of length L (or some function thereof, e.g. RMS of the charge distribution) are also scaled appropriately, and the potential is dimensionally consistent, this general property will hold.

Substituting $r = x/m$ into (2.32), m can be factored out from all the matrix elements and moved to the right side of the equation, effectively scaling the energy values

$$\begin{pmatrix} 1 + Z\alpha\Phi(x) & -\frac{\mathrm{d}}{\mathrm{d}x} + \frac{\kappa}{x} \\ \frac{\mathrm{d}}{\mathrm{d}x} + \frac{\kappa}{x} & -1 + Z\alpha\Phi(x) \end{pmatrix} \begin{pmatrix} P(x) \\ Q(x) \end{pmatrix} = \frac{E}{m} \begin{pmatrix} P(x) \\ Q(x) \end{pmatrix}$$

The eigenvalues of the equation are measured in mc^2 and the radial coordinate in $\frac{\hbar}{mc}$, where it should be stressed that m is the mass of the particle under consideration and in general different from the electron mass m_e. In other words, these quantities are no longer in natural units.

This form of the equation has multiple benefits. First, the fine-structure constant α, which is an experimentally measured value and therefore is not exact, only shows up as in the coupling of the nuclear potential. Hence, a numerical solver could simply take $Z\alpha$ instead of Z as an argument, and so the output values would no longer have any inherent experimental uncertainty attached to them. It will be explicitly up to the user to correctly convert between their physical Z value and the $Z\alpha$ parameter.

Secondly, it is no longer necessary to treat the mass of the particle m as an independent parameter. Instead, to solve the equation for a different mass value, the nuclear potential could be scaled appropriately instead. Or vice versa, the mass could be scaled when the size of the nucleus is changing. As a specific example, for a cut-off coulomb potential

$$\Phi_{\text{c.o.}}(r; R) = \begin{cases} 1/R, & 0 \leq r \leq R \\ 1/r, & r > R \end{cases}$$

it holds that $\Phi_{\text{c.o.}}(x/m; R) = m\Phi_{\text{c.o.}}(x; Rm)$. In other words, rather than changing the mass, the replacement $R \to Rm$ yields the same solution, with appropriate scalings on r and E. More generally, if scaling the parameters of the potential does not quite work, any potential could be scaled according to

$$\tilde{\Phi}(x) = \frac{1}{m}\Phi(x/m)$$

2.7 Nuclear charge distributions and potentials

So far, in the context of an external field Dirac equation (2.14) (including its radial counterpart (2.30) in Section 2.6), the potential $V(r)$ has been a generic function. And the potential, in fact, can take on many forms, depending on the application. For example, it may be an effective potential generated by the electrons close to the atomic nuclei, representing a pseudopotential [81], or exhibit an infinite potential wall when modelling confinement [82].

The most common case, though, is simply a pure Coulomb potential generated by a nuclear charge distribution. However, nuclei themselves are complicated quantum objects and there are various models that attempt to approximate the shape of a nucleus. This section reviews the various nuclear charge distribution models and their resulting potentials that are relevant in this thesis. For the various models, we follow the conventions from the very comprehensive paper by Andrae, which provides an excellent review of the different nuclear models. [83]

Nuclear potential

For the purposes of this section, we assume that the nuclear potential is generated by a spherically symmetric nuclear charge distribution $\rho(r)$ centred at the origin. For each nuclear model, the

charge distribution will be normalized to the nuclear charge Z such that

$$\int \mathrm{d}\boldsymbol{x}\ \rho(|\boldsymbol{x}|) = 4\pi \int_0^\infty \mathrm{d}r\ r^2\rho(r) = Z$$

This is convenient for the purposes of separating the shape of the nuclear potential from the particular nucleus, allowing us to easily parametrize calculations in terms of Z.

The potential $V(r)$, which is also radially symmetric if the charge distribution is, can formally be derived from the charge distribution by solving the radial Poisson equation

$$\frac{1}{r}\frac{\mathrm{d}^2}{\mathrm{d}r^2}\left[rV(r)\right] = 4\pi\rho(r) \tag{2.33}$$

with the boundary condition $\lim_{r\to\infty} V(r) = 0$. According to equation (22) in [83], the potential can be calculated by integrating the density in the following manner

$$V(r) = -\frac{4\pi}{r}\left[\int_0^r s^2\rho(s)\,\mathrm{d}s + r\int_r^\infty s\rho(s)\,\mathrm{d}s\right]$$

This expression can be obtained by using the Green's function method to solve equation (2.33).

Nuclear models

Each nuclear charge distribution model is characterized by zero or more parameters that determine its shape. It is desirable to keep the nuclear models physically comparable, which can be done by requiring the root-mean-square (RMS) radius

$$R_{\mathrm{rms}} = \sqrt{\langle r^2\rangle} = \sqrt{\frac{4\pi}{Z}\int_0^\infty \mathrm{d}r\ r^{k+2}\rho(r)}$$

of the charge distribution to be the same for all nuclear models for the same nuclear charge.

In general, when performing calculations that have a direct experimental connection, the exact value for the isotope under study would ideally also be determined empirically from experiment, such as the values tabulated in [84]. However, an experimental value is not always available, and so a very common fallback in the atomic physics community is to use the empirical fit derived Johnson and Soff in 1985 [85]: given a nuclear isotope, defined by the nuclear charge Z and nuclear mass number A, the RMS value is approximated by

$$R_{\mathrm{rms}} = (0.836A^{1/3} + 0.570) \pm 0.05\ \mathrm{fm} \tag{2.34}$$

It is worth pointing out that, due to the cube root behaviour, the RMS value of the nucleus does not change too much, qualitatively speaking. As an illustration, going from hydrogen, where $R_{\mathrm{rms}}(^1\mathrm{H}) = (1.41 \pm 0.05)\ \mathrm{fm} = (2.66 \pm 0.09) \times 10^{-5}a_0$, to oganesson, where $R_{\mathrm{rms}}(^{294}\mathrm{Og}) = (6.13 \pm 0.05)\ \mathrm{fm} = (11.58 \pm 0.09) \times 10^{-5}a_0$, the change in the RMS value, as predicted by Eq. (2.34), is less than an order of magnitude. This in turn implies that qualitatively and from a methods perspective, the results for one particular nuclear radius can be generalized for the whole range of physical nuclei.

The particular nuclear models used in this work are

- **Point-like charge distribution.** This is the standard point nucleus

$$\rho(r) = \frac{Z}{4\pi r^2}\delta(r)$$

and leads to the potential

$$V(r) = -Z/r$$

The root mean square radius here, of course, is always zero and there are no parameters to control the size or the shape of the nucleus.

- **Hollow nucleus.**[31] The defining parameter is the radius R of the shell where all the charge is located

$$\rho(r) = \frac{Z}{4\pi r^2}\delta(r - R)$$

The charge distribution leads to the potential

$$V(r) = \begin{cases} -Z/R, & 0 \le r \le R \\ -Z/r, & r > R \end{cases}$$

and root-mean-square radius $R_{\rm rms} = R$.

- **Homogeneous or uniformly charged sphere.** The defining parameter is again the radius of the sphere R, with constant charge density within the sphere

$$\rho(r) = \begin{cases} \dfrac{3Z}{4\pi R^3}, & 0 \le r \le R \\ 0, & r > R \end{cases}$$

which leads to a quadratic potential within the sphere

$$V(r) = \begin{cases} -\dfrac{Z}{2R}\left[3 - \left(\dfrac{r}{R}\right)^2\right], & 0 \le r \le R \\ -\dfrac{Z}{r}, & r > R \end{cases}$$

The root mean square radius is $R_{\rm rms} = \sqrt{3/5}R$.

- **Gaussian charge distribution.** The parameter R that we use to define the Gaussian distribution normalizes the exponent to a dimensionless number $x = r/R$ in the distribution $\rho(r) \propto \exp(-x^2)$. The precise charge distribution is given by

$$\rho(r) = \frac{Z}{R^3\sqrt{\pi^3}} \exp\left[-\left(\frac{r}{R}\right)^2\right]$$

and the related potential can be expressed via the error function $\mathrm{erf}(x)$ as

$$V(r) = -\frac{Z}{r}\mathrm{erf}\left(\frac{r}{R}\right)$$

The root mean square radius can be obtained from R with $R_{\rm rms} = \sqrt{3/2}R$.

[31] Also referred to as the "uniformly charged spherical shell" or "top-slice" model.

- **Two-parameter Fermi distribution.** This is a widely used model that depends on the two parameters a and c

$$\rho(r) = \frac{Z}{1 + \mathrm{e}^{(r-c)/a}}$$

In the limit $a \to 0$, this model becomes identical with the spherical nucleus with $R = c$. As $\rho(c) = Z/2$, c is sometimes referred to as the *half-density radius*.

A reason why this model is often preferred is because it does not have a sharp cutoff and can describe the "skin" of a nucleus — a region of finite thickness on the surface of the nucleus where the charge density falls to zero.[32]

There is no simple closed-form solution for the potential, and the RMS value is similarly complicated. However, the 1992 paper by Parpia and Mohanty describes a practical way of evaluating the potential [86]. An interesting point here is that, with two independent parameters, specifying just the RMS value is not sufficient to uniquely fix the parameters of the model. In the atomic physics community, a skin thickness $t = 2.30$ fm, chosen by Parpia and Mohanty in their paper without any particular motivation, has largely become the standard value used to fix that ambiguity.

[32] The skin thickness t, which is defined as the distance over which the nuclear charge density drops from 90% to 10% (i.e. $\rho(c - t/2) = 0.9Z$ and $\rho(c + t/2) = 0.1Z$), is related to the parameters of the two-parameter Fermi distribution via

$$t \approx 4 \ln(3) \cdot a$$

which holds well when $a \ll c$, i.e. the skin is relatively thin [83, Eq. (185) in 2.4.7, pg. 465].

Chapter 3

The Dirac equation: a numerical odyssey

As discussed in Section 2.6, the single-particle radial Dirac equation with a generic external potential poses various challenges. There are no analytical solutions available for a general potential, and the analytical approaches that do exist are tailored for particular potentials [83]. Furthermore, when treating the many-electron problems in atomic physics, for example, either in the Dirac-Hartree-Fock or a perturbative formalism, solving the single-particle equation is often a key piece of the whole approach.

This chapter explores the challenges and opportunities of numerically solving the one-particle Dirac equation for a bound state using a finite basis set. In this approach, parameters can be varied and potentials swapped out with relative ease, meaning it is very easy to study various edge cases. Although some very early prototyping was done with a B-spline basis [87], the final implementation uses the one-dimensional finite element method (FEM) basis implemented in the HelFEM program [88]. This is by no means the only numerical scheme we could have employed. For example, in atomic physics and quantum chemistry, B-splines basis sets [89, 90, 91, 92, 93], as well as ones based on exponential functions[1] [94] and Gaussians [95, 96, 97] are heavily used. Outside of basis sets, one can also simply represent orbitals on a radial grid and use a shooting method [98] or, for example, approximate a basis set with the FEM-adjacent FE-DVR[2] method [99].

The choice of FEM in particular was a pragmatic one. Unlike the other methods, the HelFEM program provided a well-structured C++ codebase implementing the basis set and the necessary functions for evaluating the operator matrix elements in that basis set. After some testing, it was also clear that the implementation and the method were very reliable and did not lead to any significant numerical difficulties. The latter contrasts with the earlier experiments with B-splines that led to various numerical issues (although it is unclear if those were due to the method or, more likely, due to the particular implementation used).

In order to systematically build up a FEM-based solver for the Dirac equation for physically interesting problems, the work proceeded in the following phases:

1. As discussed in Section 2.7, in relativistic theory the bare $1/r$-potential is often no longer good enough, as the error this approximation introduces is comparable to the relativistic correction itself, especially for the heavier nuclei. As such, it was necessary to extend HelFEM to support discretizing finite nuclear models, which brought its own set of challenges. In order to clearly separate problems arising from using realistic nuclear charge

[1] More commonly known as "Slater-type orbitals" or STOs, which have a radial part that looks like $r^n \mathrm{e}^{-\xi r}$.
[2] FE-DVR stands for "finite element, discrete variable representation".

distributions from the issues arising from the relativistic equation, this part of the work implements the non-relativistic equations.

2. The second step was to solve the standard external-potential Dirac equation for various potentials, showing that it leads to good results, but also to some challenges.

3. Finally, the same method is applied to a Dirac equation where the negative energy states have been removed via projection, showing how technically simple it is to expand this approach to more complex problems.

What resulted from this work was a Julia-language [100, 101] based solver for the Dirac equation that uses the HelFEM library underneath, used to demonstrate the performance of the FEM basis set for this problem. This high-level language implementation is easily be expanded to investigate physically and mathematically interesting edge cases of the Dirac problem, as demonstrated by the investigations of the projected Dirac equation and the diving problem.

3.1 Discretizing a differential equation

The primary task in this chapter is to determine the eigenfunctions and eigenvalues of a linear differential operator $\mathcal{L}$. For the most part, it will be the Hamiltonian H of the physical system (or some approximations of the system), and in that case the eigenstates and eigenvalues are the stationary states and their energies, respectively. As the finite element method will be applied to several different equations, it is useful to, very briefly, establish some basic, consistent notation, which in fact would work with any finite basis.

Differential operators, functions and domains

The linear operator $\mathcal{L}$ acts on a set of functions $\psi : X \to Y$ with some domain X and co-domain Y. In quantum mechanics, the co-domain is generally the set of complex numbers, but the wavefunction can have multiple components (e.g. spin components in non-relativistic quantum mechanics, or the four components of a Dirac field), and therefore the co-domain will be $Y = \mathbb{C}^m$ for an m-component function. The components ψ_α of the wavefunction are assumed to be arranged in a column vector

$$\psi = \begin{pmatrix} \psi_1 \\ \psi_2 \\ \vdots \\ \psi_m \end{pmatrix}$$

The linear operator $\mathcal{L}$ acting on a multi-component function ψ can mix the components, and so the operators will, in general, be a matrix of linear operators $\mathcal{L}_{\alpha\beta}$ acting on complex functions on the domain X

$$\mathcal{L} = \begin{pmatrix} \mathcal{L}_{11} & \mathcal{L}_{12} & \cdots & \mathcal{L}_{1m} \\ \mathcal{L}_{21} & \mathcal{L}_{22} & \cdots & \mathcal{L}_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{L}_{m1} & \mathcal{L}_{m2} & \cdots & \mathcal{L}_{mm} \end{pmatrix} \tag{3.1}$$

The exact definition of the domain X can vary substantially. The underlying physical equations generally assume the three-dimensional space as the domain (i.e. $X = \mathbb{R}^3$), but this can often be reduced. With the help of symmetries, the problem can often be solved independently and

analytically along some of the coordinates, with the numerical basis set only necessary for a subset of the coordinates (e.g. the radial coordinate in the case of an atomic system).

Each domain also comes with an inner product for calculating overlaps between quantum states, which can be written as the integral

$$\langle f_\alpha | g_\beta \rangle = \int_{x \in X} f_\alpha^*(x) g_\beta(x) w(x) \, \mathrm{d}X$$

where $w(x)$ is a weight function specific to the domain which can scale different parts of the domain differently. It is often the trivial $w(x) = 1$, but can be something non-trivial when for example coordinate transformations are involved (e.g. in spherical coordinates). In the case of multiple components, the inner product reduces to a dot product of the components multiplied by the weight function, i.e.

$$\langle f | g \rangle = \int_{x \in X} f^\dagger(x) \cdot g(x) w(x) \, \mathrm{d}X = \sum_\alpha \int_{x \in X} f_\alpha^*(x) g_\alpha(x) w(x) \, \mathrm{d}X$$

A few examples of domains X, co-domains Y, and weight functions that are relevant for this chapter are

1. *1D particle in a box.* In this case, the wavefunction is generally assumed to be a simple complex function, and the particle to be constrained in a one-dimensional interval between a and b.
$$X = [a, b], \quad Y = \mathbb{C}, \quad w(x) = 1$$
2. *Non-relativistic Schrödinger equation for electrons.* In this case, the wavefunction goes over all of space, and the wavefunction has two components — one of each of the spin-1/2 components.
$$X = \mathbb{R}^3, \quad Y = \mathbb{C}^2, \quad w(x) = 1$$
3. *Dirac equation.* The Dirac equation also has the whole three-dimensional space as its domain but has four independent components.
$$X = \mathbb{R}^3, \quad Y = \mathbb{C}^4, \quad w(x) = 1$$
4. *Radial Dirac equation.* The radial Dirac equation is an example of applying symmetries where, by going into spherical coordinates, the angular part can be solved analytically and only the function along the r coordinate is non-trivial. This also reduces the number of components down to two from four.
$$X = [0, \infty), \quad Y = \mathbb{C}^2, \quad w(r) = r^2$$
The non-trivial weight function comes from the Jacobian of the change of variables into spherical coordinates. However, it is common to define the radial functions as $\frac{1}{r}P(r)$ instead, in which case the weight function is again $w(x) = 1$, because the $\frac{1}{r}$ factor will cancel out the Jacobian.

In addition to just defining the linear operator, it is also necessary to define certain constraints on the functions ψ, such as normalizability and boundary conditions, with the latter usually requiring the function to be zero on the boundary of the domain. In all cases, they reduce the space in which the solutions are being sought.

Something else to note is that, in practice, the domain of the function may be further reduced. For example, instead of solving the problem on $r \in [0, \infty)$, it might be restricted to $r \in [0, r_{\max})$, which can be necessary when e.g. the method used can not be extended all the way to infinity with a finite number of basis functions.

Discretizing with basis sets

Discretizing an eigenvalue problem with a basis set, which is by no means the only approach, has the appeal of being conceptually very straightforward. A functional eigenvalue problem becomes a simpler matrix eigenvalue problem, which, in turn, is easy to solve using standard linear algebra techniques.[3] In principle, with a good basis set, systematically increasing the basis size should bring you closer and closer to the exact functional solution, for both the energy and the wavefunction.

In the most general notation, each component ψ_α is assumed to have a separate basis set of explicitly defined functions $\{b_i^\alpha(x)\}$ on the domain X (i.e. $b_i^\alpha : X \to \mathbb{C}$; although in practice, most often, the co-domain of the basis functions will be real numbers, so $b_i^\alpha : X \to \mathbb{R}$). Each component can then be expanded as a superposition

$$\psi_\alpha(x) = \sum_i c_i^\alpha b_i^\alpha(x)$$

where $c_i^\alpha \in \mathbb{C}$. Of course, a finite basis set can not be used to represent arbitrary functions perfectly accurately, it it is assumed that there is a systematic way to increase the size of the basis set, leading to greater and greater accuracy. While often the basis set will be the same for all components, that is not always the case. For example, when using B-splines for the relativistic Dirac equation, it is necessary to use different orders of polynomials for the large and small components to avoid numerical problems [102, 103].

In order to combine the separate basis sets into a single basis of a multi-component function, it is handy to introduce the vectors ξ^α, which are simple column vectors where the α-th component is set to unity and the rest are zero (i.e. $\xi_\beta^\alpha = \delta_{\alpha\beta}$). With their help, a multi-component function ψ is expanded as

$$\psi(x) = \sum_\alpha \sum_i c_i^\alpha b_i^\alpha(x) \xi^\alpha$$

The products $b_i^\alpha(x)\xi^\alpha$ act as multi-component basis functions and the index i can run over the different sets of basis functions for the different components.

Bra-ket notation

When the precise functional nature of the functions and operators is not important, it is convenient to use the *bra-ket notation*[4] to represent the functions, inner products etc. For example, a basis function could then just become

$$|\alpha i\rangle = b_i^\alpha(x)\xi^\alpha$$

and when it is not necessary to differentiate between the components, there could be just a single label for the basis vectors (i.e. $|i\rangle$).

Matrix elements

Given a linear operator A, made up of linear operators acting on individual components as in equation 3.1, the matrix elements are calculated as

$$A_{ij}^{\alpha\beta} = \langle \alpha i | A | \beta j \rangle = \int_{x \in X} b_i^{\alpha *}(x)\ A_{\alpha\beta}\ b_j^\beta(x)\ w(x)\, \mathrm{d}X$$

[3] Conceptually speaking, that is. In practice, the matrices often get large and it is always possible to run into numerical problems. Diagonalizing large matrices and matrices with special structures is its own field of study in applied mathematics.

[4] Explicitly proposed as a notation for quantum mechanics by Paul Dirac to have a more natural way of working with states in a Hilbert space [104]. It is hence also often referred to as the *Dirac bra-ket notation*.

While mostly $A_{ij}^{\alpha\beta}$ should be interpreted as elements of a matrix with indices (α, i) and (β, j), it is also useful to remember its block matrix structure

$$A = \begin{pmatrix} A^{11} & A^{12} & \cdots & A^{1m} \\ A^{21} & A^{22} & \cdots & A^{2m} \\ \vdots & \vdots & \ddots & \vdots \\ A^{m1} & A^{m2} & \cdots & A^{mm} \end{pmatrix}$$

where $A^{\alpha\beta}$ are block matrices corresponding to different components of the multi-component function.

Working in non-orthonormal bases

In practice, most basis sets used to represent functions, whether in FEM or otherwise, are not orthonormal, meaning that some operations require a bit more care. Since quantum theory is mostly taught assuming orthonormal bases, it is not always obvious where this might become an issue.

If V is a vector space and a finite, linearly independent (but not necessarily orthonormal) basis set $\{|i\rangle\}$ spans it or a subspace of it, then any vector will have a unique expansion

$$|\Psi\rangle = \sum_i c_i |i\rangle$$

with coefficients c_i. Similarly, linear operators $\hat{T} : V \to V$ can be represented using a matrix of coefficients

$$\hat{T} = \sum_{ij} T_{ij} |i\rangle \langle j|$$

Now, in an *orthonormal* basis, the expansion coefficients for both vectors and operators can be calculated by employing the inner product

$$c_i = \langle i|c\rangle , \quad T_{ij} = \langle i|\hat{T}|j\rangle$$

The expression $\langle i|\hat{T}|j\rangle$ is the matrix element of the operator between the basis states $|i\rangle$ and $|j\rangle$. However, in a non-orthonormal basis, things are not quite as simple. In this case, the overlap matrix S, defined as

$$S_{ij} = \langle i|j\rangle \neq \delta_{ij}$$

is not the identity matrix and therefore the explicitly calculated overlaps and matrix elements $\tilde{c}_i \equiv \langle i|c\rangle$ and $\tilde{T}_{ij} \equiv \langle i|\hat{T}|j\rangle$ are not, in general, equal to the expansion coefficients. Instead, they can be related by using the overlap matrix S. For the vector expansion coefficients

$$\tilde{c}_i \equiv \langle i|c\rangle = \sum_j \langle i|j\rangle c_j = \sum_j S_{ij} c_j = (Sc)_i \neq c_i$$

and for the matrix elements

$$\tilde{T}_{ij} \equiv \langle i|\hat{T}|j\rangle = \sum_{kl} S_{ij} c_j = (Sc)_i \neq c_i$$

As the overlap matrix S is just a matrix of different inner products, it is positive-definite and therefore invertible. With that in mind, we can always invert the previously derived relations

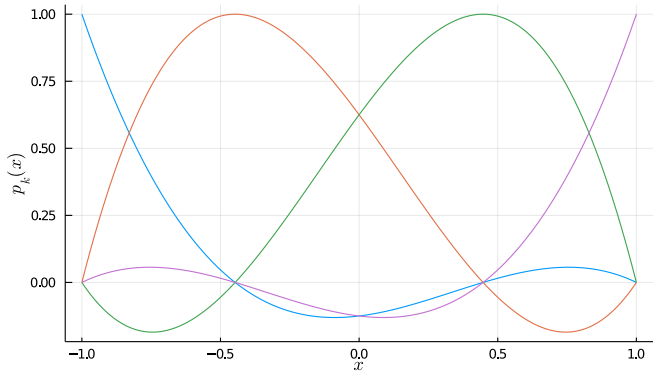

Figure 3.1: Illustrative example of a polynomial basis used within each element in the FEM method. In this case, these are the four Lagrange interpolating polynomials defined by four Gauss-Lobatto nodes on the domain $x \in [-1, 1]$.

and recover the inherent expansion coefficients from the explicitly calculated matrix elements and overlaps

$$\tilde{c} = Sc \iff c = S^{-1}\tilde{c}, \qquad \tilde{T} = STS \iff T = S^{-1}\tilde{T}S^{-1}$$

These relationships do not have any real physical significance but are important when implementing methods, e.g. to make sure that functions are re-constructed correctly.

3.2 Finite Element Method

The finite element method (FEM), while better known in the context of multivariable partial differential equations and its applications in engineering [105], can also be used to elegantly solve one-dimensional problems and applied to quantum theory [106]. It was picked for this work as it has been shown to be reliable and performant for atomic systems, and there exists a well-written implementation in the form of the HelFEM software [88].

In the one-dimensional FEM the domain $X = [a, b]$ is divided up into sub-intervals (elements; with element boundaries $r_0 = a < r_1 < \cdots < r_N = b$ for N elements). The idea is not to cover the whole domain with increasingly complicated functions, but instead rely on a set of simple functions — the local basis polynomials $\{p_k(x)\}$ — on some sub-interval and repeat them until the whole domain is covered. To form the basis functions $\{b_i(r)\}$, those local basis functions $p_k(x)$, conventionally defined on the domain $x \in [-1, 1]$, are essentially just repeated in each element by appropriately scaling the argument of the function to change the domain from $[-1, 1]$ to $[r_i, r_{i+1}]$. To impose boundary conditions, some functions are removed from the first and the last elements, and also functions that have non-zero values on the element boundary are joined into a single basis function spanning two elements. In principle, the intervals, defined by a grid of interval boundaries on the domain, can be chosen arbitrarily, and by making the grid finer and finer, it should be possible to achieve convergence and a better representation of the solutions of the differential equations. However, a good choice of the grid significantly improves the rate of convergence.

While HelEFM implements several options for the local functions, this work exclusively uses the Lagrange interpolating polynomials (LIPs) with Gauss-Lobatto nodes, as those were determined to be the best-performing ones for atomic systems [88]. The exact definitions for the LIPs

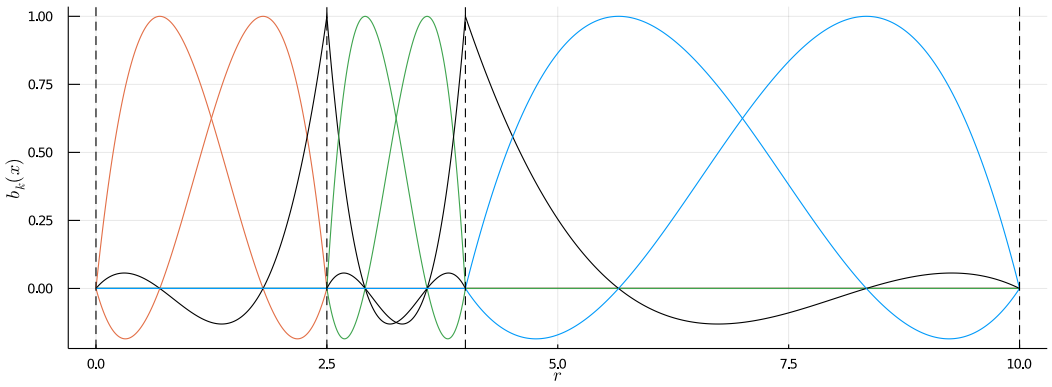

Figure 3.2: An example of a FEM basis with three unevenly distributed elements on the domain $r \in [0, 10]$. The vertical dashed lines represent element boundaries. The polynomials within each elements are the same as in Figure 3.1, with the x-axis linearly scaled as appropriate. The basis functions on the left-most and right-most boundary have been removed to impose boundary conditions $f(0) = 0$ and $f(10) = 0$. Note that the basis functions have non-zero values only in each element (illustrated using the colors), except the ones that have non-zero values on the boundaries (black lines) — these ones are constructed by joining the edge-most functions in each element together to form a single basis function that spans two elements.

can be found in the literature, Figure 3.1 illustrates what they generally look like. An illustration of a fully constructed one-dimensional FEM basis can be seen in Figure 3.2.

As all the basis functions are all only defined on a finite subsection, the basis can not, with a finite number of basis functions, be extended all the way to infinity. Therefore it is necessary to pick a finite domain for the calculations, which in the case of radial atomic problems means solving the equations on $r \in [0, r_{\max}]$ (corresponding to a three-dimensional ball), instead of the full domain $r \in [0, \infty)$. This can be interpreted as an artificial, non-physical boundary condition $\Psi(r_{\max}) = 0$ at the outer boundary and leads to a form of confinement [107]. However, for bound states, if $r_{\max}$ is large enough, this will not affect the energies and the states, since the wavefunctions effectively become zero far from the origin due to an exponential decay at large r. But it does introduce another parameter that needs to be checked for convergence when performing such calculations.

Matrix elements

The purpose of defining a basis set is to calculate the matrix elements of arbitrary linear operators in this basis. For an operator A acting on functions on the domain $X = [0, r_{\max}]$ of radial functions, the matrix elements are given by the integral

$$A_{ij} = \int_0^{r_{\max}} \mathrm{d}r \; b_i^*(r) A b_j(r)$$

HelFEM provides routines for two generic types of operators. First, it is possible to calculate the matrix element of an arbitrary function $f(r)$

$$\langle b_i | f(r) | b_j \rangle = \int_0^{r_{\max}} \mathrm{d}r \; b_i^*(r) f(r) b_j(r) \tag{3.2}$$

Secondly, it provides a routine for evaluating the matrix elements of r^k, possibly also by calculating it with the first derivatives of the basis functions

$$\langle \partial^m b_i | r^k | \partial^n b_j \rangle = \int_0^{r_{\max}} \mathrm{d}r \; \frac{\partial^m b_i^*(r)}{\partial r^m} r^k \frac{\partial^n b_j(r)}{\partial r^n} \tag{3.3}$$

where $m, n \in \{0, 1\}$ and $k \in \mathbb{Z}$.

While HelFEM does not provide an explicit routine for evaluating the matrix elements of the second derivative, the fact that the basis functions have non-zero values only on a finite interval means that it can be evaluated from the first derivatives. By noting that for any pair of basis functions $b_i(r)$ and $b_j(r)$ there exists an interval $[r_p, r_q]$ outside of which the product of the basis functions is zero (i.e. $b_i(r_p) = 0 = \phi_i(r_q)$), it can be shown with the help of partial intergration that

$$\begin{aligned}\langle b_i | \partial^2 | b_j \rangle &= \int_0^{r_{\max}} \mathrm{d}r \; b_i^*(r) \frac{\partial^2}{\partial r^2} b_j(r) \\ &= b_i^*(r) \frac{\partial b_j(r)}{\partial r}\bigg|_0^{r_{\max}} - \int_0^{r_{\max}} \mathrm{d}r \; \frac{\partial b_i^*(r)}{\partial r} \frac{\partial b_j(r)}{\partial r} = - \langle \partial b_i | \partial b_j \rangle\end{aligned}$$

In other words, evaluating the matrix elements of the second derivative operator reduces to appropriately evaluating (3.3).

Calculating the matrix elements of the Hamiltonian of the system is simply a matter of adding together the matrix elements of all the basic building block operators, with (3.2) and (3.3) covering everything needed in most atomic physics calculations. The eigenvalue problem to determine the stationary states and eigenenergies of the system can then be solved by feeding the Hamiltonian matrix to any matrix diagonaliser. However, what needs to be kept in mind is that the basis is not orthonormal, so one has to solve a generalized eigenvalue problem

$$H\phi = ES\phi$$

where S is the overlap matrix $S_{ij} = \int_0^{r_{\max}} \mathrm{d}r \; b_i^*(r) b_j(r)$. While linear algebra libraries usually do have generalized eigenvalue solvers available, here the basis was first effectively orthogonalizing with Löwdin orthogonalization [108], consistent with the original HelFEM paper. That is, the Hamiltonian matrix was transformed as

$$H' = S^{1/2} H S^{-1/2}$$

and it was the matrix H', representing the operator in an orthonormal basis, that was passed on to a standard matrix diagonaliser. To reconstruct the state vectors ϕ in the original basis, the eigenstates ϕ' or H' can be rotated back the original non-orthonormal basis by the matrix multiplication

$$\phi = S^{-1/2} \phi'$$

Whether using a generalized eigenvalue solver or the Löwdin orthogonalization, it obviously does not lead to an exact result, as the eigenvalues are determined using a finite basis set. However, the eigenvalues and eigenfunctions are expected to converge to the exact values in the limits of increasing either the polynomial order or by dividing the elements further and in the limit $r_{\max} \to \infty$.

Numerical quadrature

In HelFEM, the integrals (3.2) and (3.3) are evaluated numerically with the Gauss-Chebyshev quadrature. That is, the matrix elements A_{ij} of the operator A are approximated by summing up the values N_{quad} points within an element located at r_n

$$A_{ij} = \int_0^{r_{\max}} \mathrm{d}r\; b_i^*(r) A b_j(r) \approx \sum_{n=1}^{N_{\mathrm{quad}}} w_n (b_i^\dagger \hat{A} b_j)(r_n)$$

with each associated with a weight w_n. This offers more numerical stability than the analytical formulae, mostly due to the numerical instabilities introduced by the scaling, which happens when the polynomials are translated to the elements, and it makes is trivial to implement integrals of arbitrary functions. However, it does introduce another parameter — the number of quadrature points — that needs to be chosen carefully. It is important to have enough quadrature points for the integrals to be accurate, but adding extra points will have an associated computational cost.

3.3 Finite nuclear charge distributions and FEM

This section explores the challenges of incorporating finite nuclear models into finite element calculations, with the aim to converge the energies and wavefunctions to floating point precision[5]. Even though finite nuclear corrections are usually not relevant in non-relativistic calculations, the benchmarking calculations here were performed with the non-relativistic atomic Hamiltonian in order to avoid any technical hurdles the Dirac equation might introduce.

Physical model

Mathematically, the goal is to determine the eigenstates $\Psi(\boldsymbol{x})$ and eigenvalues E of the non-relativistic time-independent single-particle Schrödinger equation in a central potential $V(r)$, which is SI units looks like

$$\left[-\frac{\hbar^2}{2m_{\mathrm{e}}}\nabla^2 + V(|\boldsymbol{x}|)\right]\Psi(\boldsymbol{x}) = E\Psi(\boldsymbol{x})$$

where $\boldsymbol{x} \in \mathbb{R}^3$ and m_{e} is the mass of the electron. The implementation of the equation is in natural units, so $\hbar = m_{\mathrm{e}} = 1$. As the potential $V(r)$, being generated from a spherically symmetric charge distribution such as the ones described in Section 2.7, is spherically symmetric, the equation reduces to a one-dimensional radial equation in spherical coordinates. A convenient way to represent the wavefunctions is via the radial function $P(r)$ defined by

$$\Psi(x) = \frac{1}{r}P(r)Y_{\ell m}(\theta,\varphi)$$

where $Y_{\ell m}$ are the spherical harmonics (see also Appendix A), corresponding to the azimuthal and magnetic angular momentum quantum numbers ℓ and m. The corresponding eigenvalue equation for $P(r)$ is

$$\left[-\frac{1}{2}\frac{\mathrm{d}^2}{\mathrm{d}r^2} + \frac{\ell(\ell+1)}{2r^2} + V(r)\right]P(r) = EP(r)$$

[5] All calculation were performed with 64-bit floating point numbers,

and must be solved for each angular momentum subspace separately.

As options for the potential $V(r)$, HelFEM currently has an implementation for four of the nuclear models discussed in Section 2.7. Each of the implemented potentials was used in the benchmark calculations, and each served a particular purpose:

1. $1/r$ **potential of a point nucleus:** As this potential was already implemented and tested for [88], it serves as a known working reference case for the new code, including the relativistic implementation. However, it is important to keep in mind for this potential that it is singular at the origin.

2. **Hollow nucleus of radius R:** This is the simplest potential that would describe a nucleus that is not a point and does not have a singularity at the origin. It is possible to determine the eigenvalues almost analytically, only requiring a numerical root-finding procedure [83], which makes this model excellent for testing the implementation for nuclei with finite extent. However, it does exhibit a point at the cutoff radius where higher derivatives do not exist (a type of singular point), which, as we shall see, will require special treatment.

3. **Homogeneous uniformly charged sphere with radius R:** This potential has a slightly smoother shape than the hollow nucleus but still exhibits a singular point at the point that corresponds to the surface of the nucleus. It allows us to double check our results for the hollow nucleus, including how we deal with the singular point in the FEM framework.

4. **Gaussian nucleus:** This potential is smooth at all points and, therefore, does not require special treatment at any point. It is also very common in relativistic Gaussian basis set based quantum chemistry and atomic calculations. As such, it would be the perfect test case for a finite nuclear model. However, unfortunately, there is no easy way calculate reference eigenvalues for this potential, and so we need the other singular potential to test the implementation before using this potential.

The two-parameter Fermi nuclear model is currently not implemented, as evaluating the potential is somewhat non-trivial. As the potential is extensively used in atomic physics when modelling nuclei in relativistic calculations, it would be useful to make it available to users of HelFEM and have it benchmarked as part of future work. However, as the model qualitatively lies between a Gaussian and spherical nucleus, the results and insights from this sections also carry over to the Fermi nuclear model.

Numerical quadrature

As was mentioned in Section 3.2, HelFEM uses numerical quadrature to evaluate the matrix elements of operators. This means that the matrix elements are approximate but converge with the increasing number of quadrature points N_{quad}. As the goal is to achieve very high accuracy, ideally converged to machine precision, it is important to make sure that the matrix elements are evaluated with enough quadrature points such that they themselves are converged.

If $(A_k)_{ij}$ is a matrix element of the operator evaluated with k quadrature points in each element, then the approximate error from the converged result can numerically be calculated with

$$\epsilon_k = \max_{i,j} \left| (A_k)_{ij} - (A_N)_{ij} \right|$$

where N is the case with the largest number of quadrature points that was calculated. In all the following calculations, the chosen value was $N = 500$ which is well within the converged region where floating-point error dominates for all the cases that were considered. This is the *maximum*

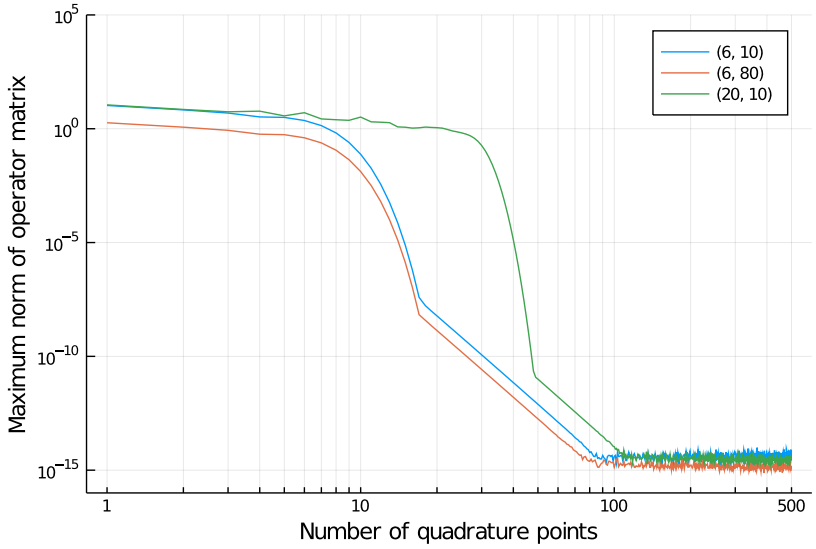

Figure 3.3: Convergence of the overlap matrices when increasing the number of numerical quadrature points, with the y-axis showing the maximum norm of the difference from the case with the highest number of quadrature points (500). Each line corresponds to a different radial basis set, defined by the (order of polynomial, number of elements) tuple. A simple exponential grid with $z = 2$ and $r_{\max} = 40$ is used for the element boundaries. Naturally, the calculation was performed only at integer x values, and the non-integer cases are simply interpolated.

norm of the matrix difference and is useful because it is the most extreme of the matrix norms — a convergence problem in a single matrix element will already stand out.

Figure 3.3 shows, on a log-log scale and for a few different basis sets, the convergence of the overlap matrix, which can also be thought of as the simplest possible operator: the identity operator. Qualitatively, the figure shows a systematic convergence, where starting from a certain number of quadrature points (different for different the polynomial orders), the error enters a clear power-law region, after which it hits the floating-point precision at around 10^{-15}. The convergence is slightly slower the higher the order of polynomials is, which is expected, since higher-order polynomials mean finer details in the basis functions. Changing the number of elements does not have a substantive effect on convergence.

Figure 3.4 demonstrates the convergence behaviour of the $1/r$ potential and kinetic energy operators compared to the overlap matrix. The $1/r$ operator converges essentially as fast as the overlap matrix. The kinetic energy, which is calculated as the integral of the derivatives of the basis elements, shows the same convergence behaviour as the other operators, but the resulting numerical noise is a few orders of magnitude larger.

Figure 3.5 shows the convergence of the operators corresponding to the finite nuclear models. What is immediately clear is that the integrals for nuclear models with non-smooth potentials converge very slowly, which is because the singular point of the potential usually falls in the middle of an element. The matrix elements for the Gaussian nuclear model do not converge quite as fast as the other operators either, although it is not as pathological as the potentials with singularities. This is due to the relatively few (10) elements used in this benchmark calculation, and therefore the first element boundary with this particular grid will be relatively far from the non-trivial region of the potential (at about $10^{-2}\ a_0$ vs the RMS radius of $10^{-4}\ a_0$). Increasing the

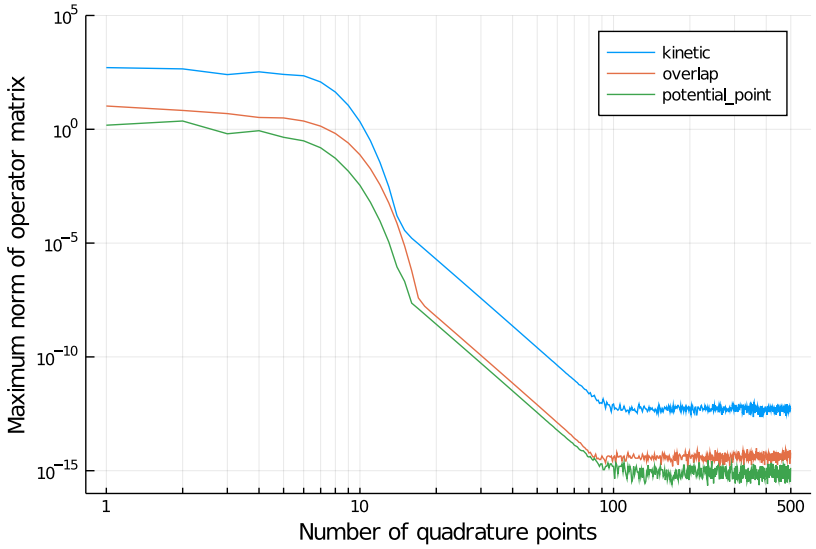

Figure 3.4: Convergence of the overlap matrix and the kinetic and $1/r$ potential operators as a function of the number of quadrature points. Axes as in Figure 3.3. The basis set is generated from an expotential grid with 10 elements and using 6th order polynomials for each element. Like in Figure 3.3, the calculation was performed only at integer x values, and the non-integer cases are simply interpolated.

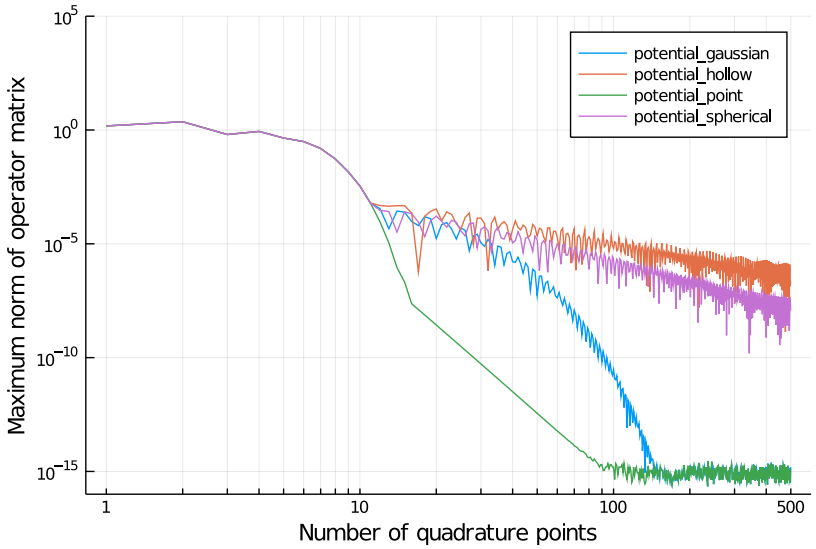

Figure 3.5: Convergence of the potential operators for the different nuclear models under consideration (point, hollow, spherical, Gaussian). Axes as in Figure 3.3. The basis set is generated from an exponential grid with 10 elements and using 6th order polynomials for each element, with no extra boundaries. Like in Figure 3.3, the calculation was performed only at integer x values, and the non-integer cases are simply interpolated.

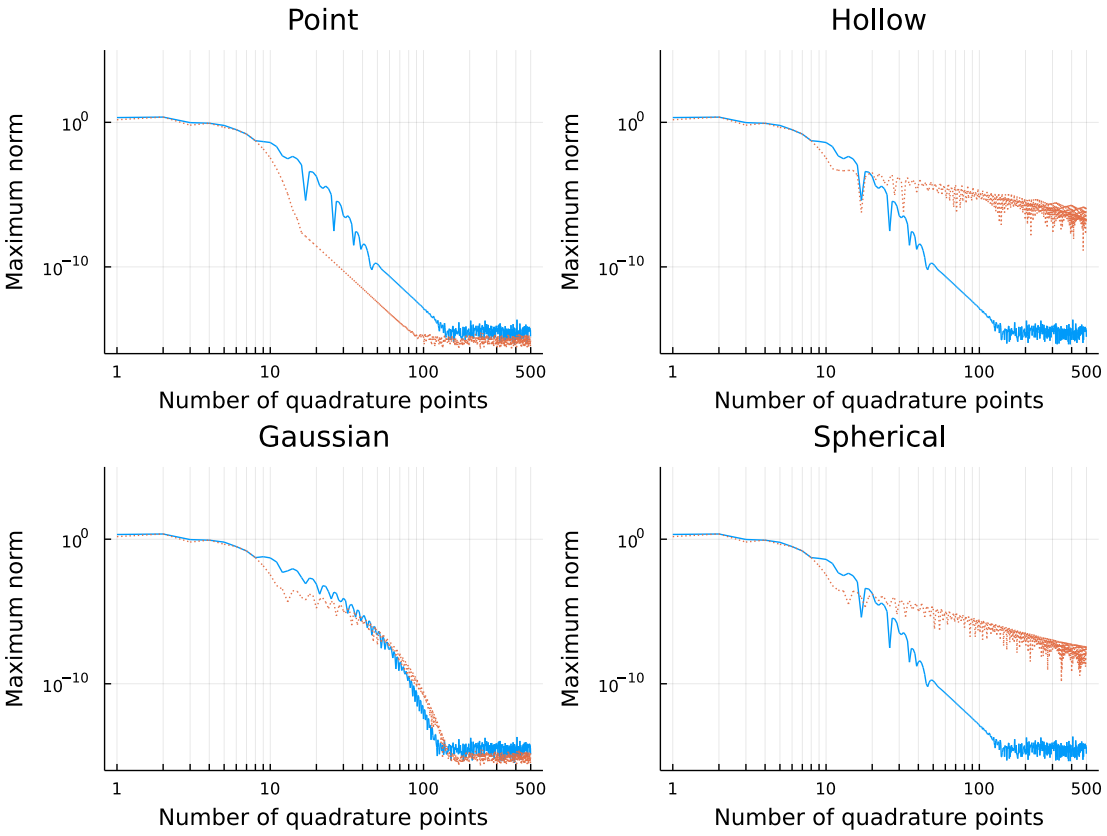

Figure 3.6: Comparison of the convergence of the potential operators for the different nuclear models depending on whether the basis has an extra boundary (solid blue lines) or not (dotted red lines). Axes as in Figure 3.3 and the basis set is generated from an exponential grid with 10 or 11 elements and using 6th order polynomials for each element. Like in Figure 3.3, the calculation was performed only at integer x values, and the non-integer cases are simply interpolated.

number of elements brings the Gaussian nucleus in line with the other nuclear models, whereas it has no effect on the convergence of the potentials with singularities, hinting that having extra elements close to the nucleus is probably helpful when working with finite nuclear models.

The convergence problems with the non-smooth potentials can be resolved by inserting an extra element boundary at exactly the singular point, splitting the problematic element in two. The effect of the extra boundary on the convergence can be seen in Figure 3.6 which compares the two such radial grids — one with an extra boundary and one without. An extra boundary at the singular point indeed fixed the convergence behaviour of the non-smooth potentials, whereas the smooth potentials are essentially unaffected by the extra boundary.

Something else that should be noted is that in the original HelFEM implementation, the default value for the number of quadrature points is five times the order of the polynomial. However, this benchmarking shows that this does not lead to machine precision, and, especially for finite nuclear models, more quadrature points are needed. Based on the benchmarks, 200 quadrature points seems sufficient to guarantee convergence to machine precision (if extra boundaries are introduced for non-smooth potentials). Alternatively, it might also be worth investigating dy-

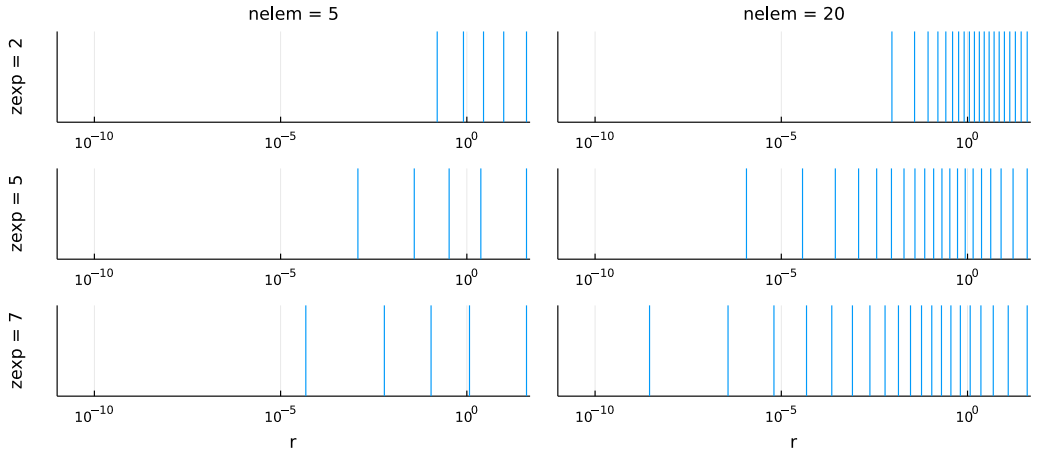

Figure 3.7: An illustration of exponential grid boundaries distributed according to Eq. (3.4) on a logarithmic r axis with increasing z parameter (vertically; $z = 2, 5, 7$) for $N = 5$ (left) and $N = 20$ (right) elements. In a FEM basis, there is a full set of polynomials between each consecutive pair of boundaries, as explained in Section 3.2. For these examples, $r_{\max} = 40\ a_0$.

namically determining the necessary number of quadrature points, or minimally at least running a simple convergence check to ensure that the operators are converged to some pre-determined accuracy.

Choice of grid

In principle, simply increasing the number of elements or the order of polynomials should lead to convergence, independent of how the elements themselves are distributed. However, it is also important to have an implementation that is efficient, i.e. having the same level of accuracy but with the lowest number of elements. For this reason, the choice of grid boundaries should also be double-checked.

As the goal is to study finite nuclear models, the most significant part of the radial axis is the part close to the nucleus. Out of the different grid choices implemented in HelFEM, it is, therefore, the *exponential grid* that should be the most suitable and efficient grid as it already concentrates elements close to the origin. Such a grid also describes more accurately the strong oscillations in the core region for higher principal quantum number wavefunctions and was shown to be the most efficient grid for non-relativistic multi-electron calculations with a point nucleus in the original HelFEM paper [88].

An exponential grid of N elements is defined by the $N + 1$ element boundaries

$$r_i = (1 + r_{\max})^{q(i)} - 1, \quad q(i) = \left(\frac{i}{N}\right)^z, \quad i = 0, \ldots, N \tag{3.4}$$

where $r_{\max}$ and z are parameters of the grid.[6] Whereas $r_{\max}$ simply determines the last, outer elements boundary ($r_N = r_{\max}$), the z parameter controls how the boundaries are distributed within the interval — the larger the z value is, the closer the element boundaries are to the

[6]The notation here for the parameters is slightly different when compared to the original paper in order to align more closely with the HelFEM C++ code.

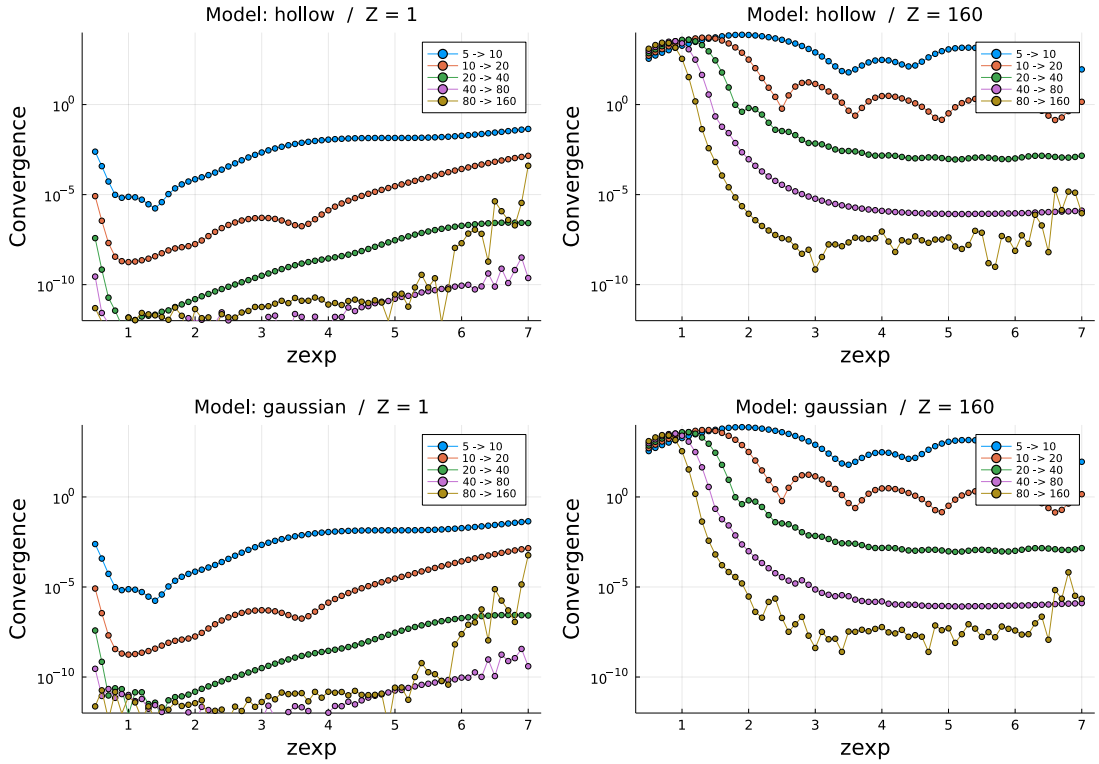

Figure 3.8: Convergence of the $1s$ energy as a function of the z parameter of the exponential grid for the hollow (top row) and Gaussian (bottom row) nuclei with $Z = 1$ (left) and $Z = 160$ (right). The calculations were performed with order 6 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\mathrm{rms}} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

origin. A few example exponential grids are shown in Figure 3.7, also illustrating the effect of changing the z parameter of the grid.

Replicating the scans of the z parameter from the HelFEM paper [88], presented in Figure 3.8 (with more figures in Appendix C), was the easiest option to assess whether the value of $z = 2.0$ recommended in the paper needs to be updated. To visualize convergence without referring to independent (e.g. analytical) reference values, the figures use the change in energy as the number of elements is doubled as the quantitative estimate for the convergence speed. A smaller value for a particular basis is assumed to indicate faster convergence, as the basis is likely closer to being converged if, for example, going from 80 to 160 elements leads to a smaller improvement. By visually inspecting the figures, it is clear that, especially at higher Z values, a z parameter larger than the recommended $z = 2$ is desired, if the goal is to have something that would universally work at both high and low nuclear charges. Recommending $z = 3$ for finite nuclei seems reasonable, and this will be used in all the following calculations with finite nuclei.

Convergence

As mentioned, the numerical approach involves a whole set of parameters that define the basis set, all of which should be checked to make sure that the eigenstates and eigenvalues have converged

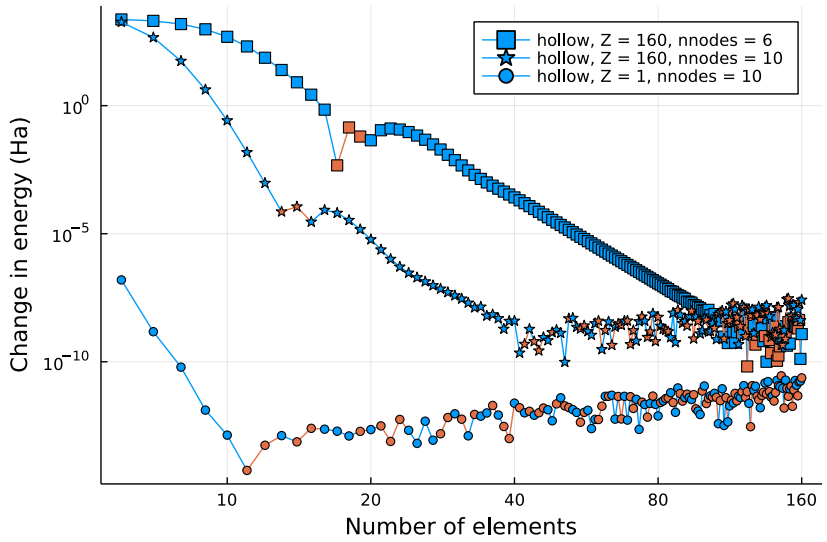

Figure 3.9: Convergence of the non-relativistic $1s$ energy with increasing number of elements for a few basis sets and nuclear charges Z with the hollow nuclear model. The y-axis shows the change in energy as another element is added (with the domain kept the same). The colours indicate whether the change in energy was positive or negative (red and blue, respectively).

to the true solution. However, verifying convergence is tricky as there are no known values to compare against necessarily. Furthermore, any rules of thumb for choosing the basis parameters that could be devised will only be valid for some range of physical parameters and it can be hard to know when they become invalid. As such, it was necessary to develop a generic procedure to reliably calculate fully converged energies.

Figure 3.9 illustrates how the energy converges as the number of elements is increased. Similar to the previous subsection, as there are no reliable reference values to compare to, the y-axes show the change in the energy between two points on the x-axis, rather than the absolute energy. There is a clear power-law behaviour visible at first as the value gets smaller and smaller, consistent with the variational principle. At one point, the machine precision is reached and the value starts oscillating (which can be seen by the changing colours) around some constant value, assumed to be the converged energy. As expected, using higher-order polynomials leads to faster convergence.

The numerical noise for greater nuclear charges appears to be larger, but this can be explained by noting that the absolute value of the energy also scales with Z: the scaling of the energy is roughly Z^2, implying a 10^4 difference in energies, which is broadly consistent with the figure. As a final observation, the numerical noise in the converged region is smallest right at the point where numerical noise takes over and then increases with the number of elements.

Algorithmic estimation of the converged energy

Having a qualitative feeling for the behaviour of the energy as it converges with an increasing number of FEM elements, it is possible to devise a simple algorithm to reliably estimate the converged energy.

Conceptually, the idea is to plot the energy as a function of the number of elements and then run a sliding fixed-width window across the x-axis, calculating an estimate for the energy in each

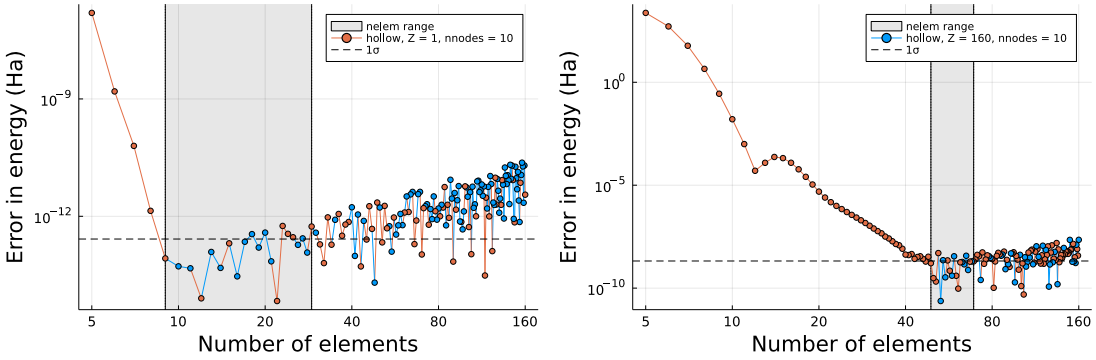

Figure 3.10: Convergence of the energy for high and low Z values ($Z = 1$ and $Z = 160$, respectively), as a function of the number of elements. On the y-axis we plot the difference from the converged estimate, which is calculated as the mean of the energy values in the gray region, with color indicating if the particular energy is lower (blue) or higher (red) than the estimate. The range of 20 element values used for the estimate, illustrated using the gray span, is determined using the sliding minimum variance approach as described in the text.

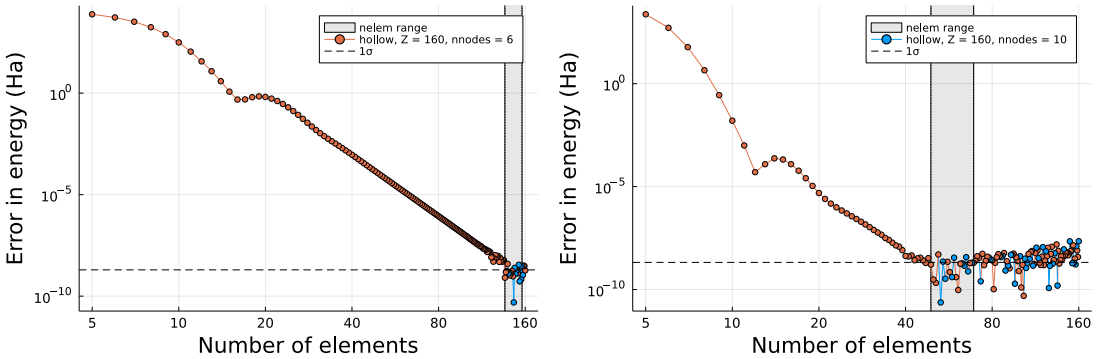

Figure 3.11: Same as Figure 3.10, but showing the difference in convergence when employing different orders of polynomials (6 and 10, respectively).

window via an average. The final estimate can then be picked by choosing the window that has the lowest variance. This works because the regions where the energy is still converging, the change in energy is large (relative to numerical precision) and mostly in one direction, which both will push up the variance. In the converged region, the energy will just oscillate around the converged value, but the amplitude of the numerical noise increases with more elements. The region with the lowest variance minimizes the numerical error of the estimate.

The width of the window was chosen to be 20 points on the x-axis, based on what seemed to work best after some trial and error. Using too few points means running the risk of accidentally picking spurious "stable" regions (i.e. where the energy happens to not change much due to random chance), whereas too many points mean the accuracy starts dropping due to the increasing numerical noise from including too many elements.

Figure 3.10 shows the algorithm in action. Like in Figure 3.9, the convergence for lower Z values is faster than for the higher ones. $Z = 1$ needs only about 10 elements to converge to

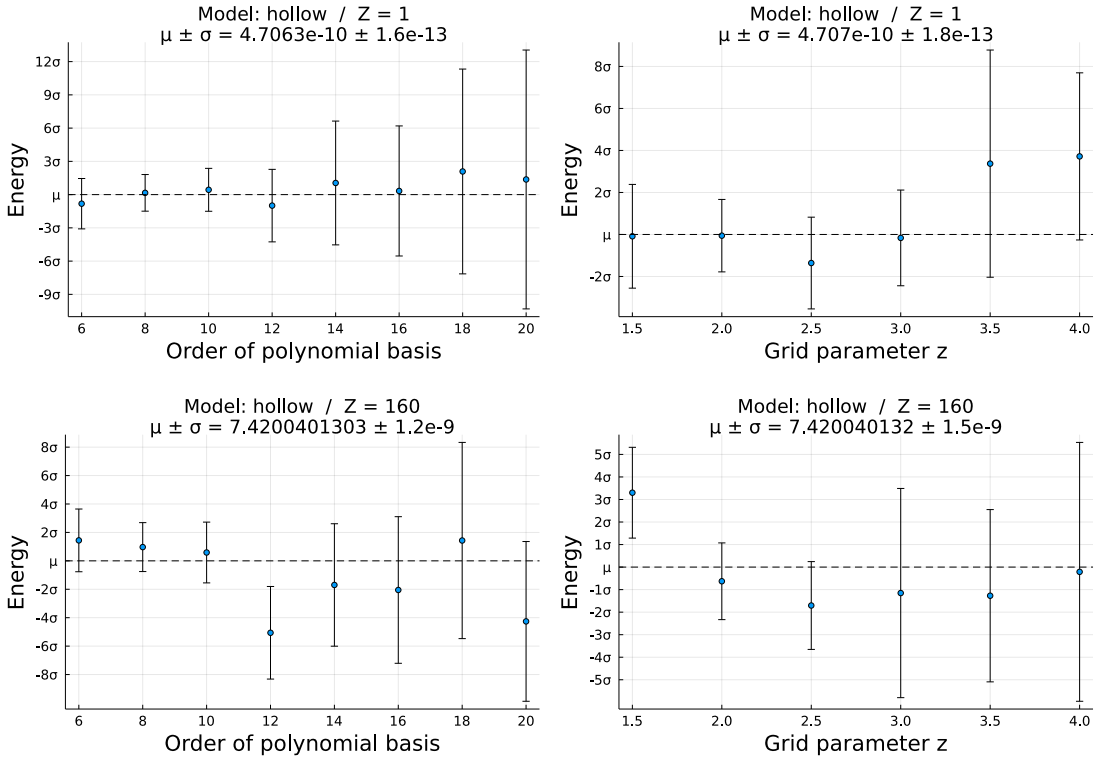

Figure 3.12: Energy estimates of the hydrogenic $1s$ energy with different basis sets for the hollow nuclear model ($R_{\text{rms}} = 10^{-4}\ a_0$) for two different Z values, illustrating the consistency of converged energy estimates. To sample different basis sets, each data point is determined with a slightly different grid of elements, with the order of polynomials varied for the values in the figures on the left (and fixed $z = 2.0$), and the z parameter of the exponential grid varied for the figures on the right (with polynomial order fixed to 10). The dashed line is the inverse-variance weighted arithmetic mean of the data points and the ticks on the y-axis are in the units of standard deviation of the weighted mean, with the exact values for both given in the figure titles.

numerical precision, whereas $Z = 160$ requires about 50–60. Either way, the algorithm finds the most optimal region to estimate the energy from. Figure 3.11 compares different orders of polynomials — the lower the order, the slower the convergence. It also shows that with 6th order polynomials, it converges very close to the maximum number of elements chosen for the calculation, which is a problematic edge case that one needs to be mindful of as the algorithm is not able to detect it automatically.

As an improvement to this procedure, it might be worth trying to quantify whether the energy in the chosen window is indeed oscillating around a converged value or is still converging to something. The correlation of energy vs. the number of elements could potentially be used to quantify this — the converged region should exhibit near-zero correlation.

While the speed of convergence may vary between different basis sets (e.g. when changing the distribution of elements or the order of the polynomials), they should all converge to the same energy in the end. This is demonstrated by Figures 3.12 and 3.13, where the algorithm gives stable results across basis a range of basis parameters.

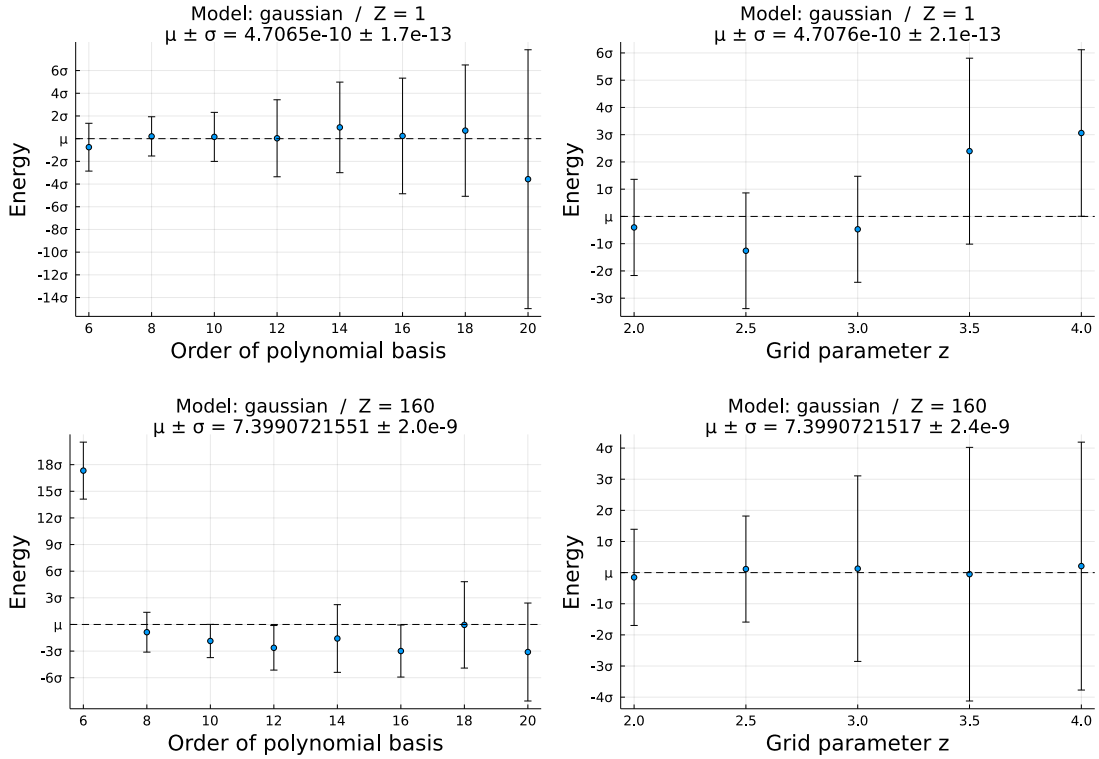

Figure 3.13: Same as Figure 3.12, but for the Gaussian nuclear model. Compared to Figure 3.12, it is not showing the $z = 1.5$ case, because it actually has not converged.

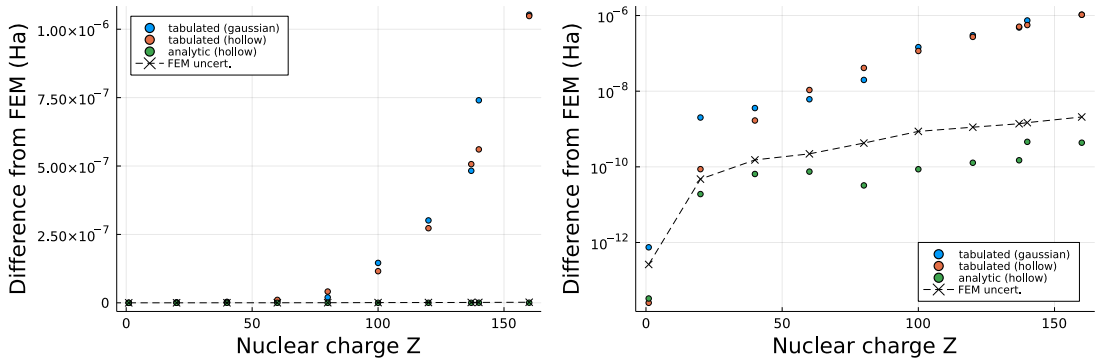

Figure 3.14: Disagreement of the tabulated energies in the Andrae reference [83] and the energies calculated with FEM (linear and logarithmic scales on the left and right, respectively). The y-axis is the energy difference from the converged FEM results. The label "tabulated" refers to the values presented in Table 4 of the reference, whereas "analytic" refers to solving the equations analytically with the approach described in the paper.

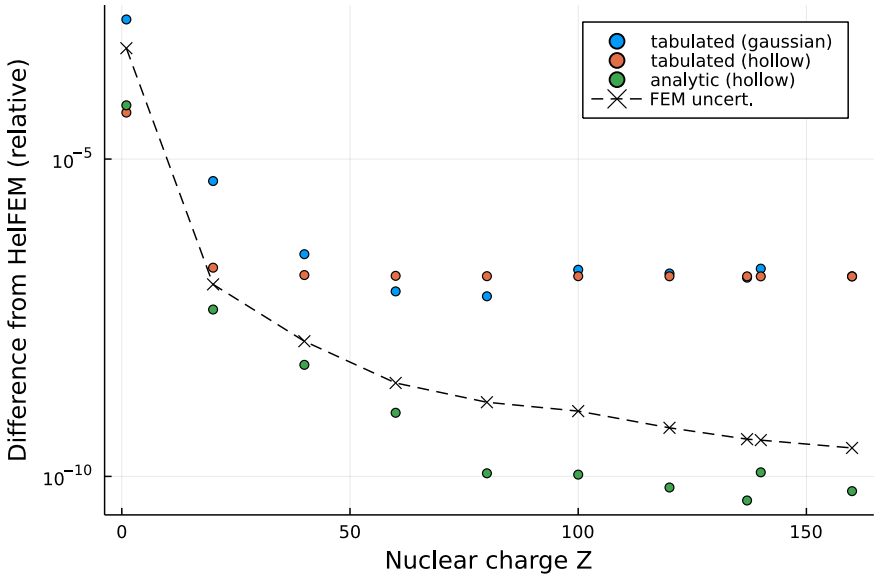

Figure 3.15: Same as Figure 3.14, but with a values relative to the estimated value of the finite nuclear correction.

Single particle benchmarks

With the algorithm in hand, it is now possible to benchmark our approach against reference values tabulated in the literature [83]. However, as can be seen on Figures 3.14 and 3.15, the values tabulated by Andrae disagree with the ones obtained from the FEM calculations. While the disagreement is very small relative to the already small finite nucleus correction (on the order of $10^{-6} - 10^{-4}$ as a fraction of the finite nucleus correction), it is systematic and much larger than the numerical noise from the FEM approach. The initial suspects were (1) non-convergence with respect to the finite element basis, or (2) non-convergence with respect to the artificial outer boundary $r_{\max}$. These were quickly ruled out — the former by the algorithm itself and the latter by manual verification[7].

Having established that the FEM results appear to be consistent, precise and correct, and assuming that the same applies for the reference values from the literature, it appears that there is an inconsistency in the setup of the physical problem being solved. The only parameter that is not clearly defined in the reference is the RMS value, which in turn defines the parameters of the nuclear models, with the paper only stating that it is obtained from Equation 2.34. It is possible that in the calculations for the reference paper the RMS value was, for example, rounded, whereas the FEM implementation uses this equation directly. The fact that the shift in energy appears to be consistent between the different nuclear models further substantiates this hypothesis.

To check this, we re-implemented the analytic matching approach described and used for the values in [83], which should yield exact results without any numerical parameters. As can be seen in Figures 3.14 and 3.15, the re-evaluated analytical results for the hollow nuclear model match the FEM results very well (within numerical noise). This shows that it is possible to achieve

[7]It should be noted here that $r_{\max} = 40$ a.u. is sufficiently large for $n = 1$ states for all Z, but at low Z values and for higher n values a larger value must be used, since the higher n orbitals do extend farther out.

Table 3.1: Energies of the hollow nuclear model as calculated with FEM and compared with the *tabulated values* from [83]. All energies given in Hartrees.

Z	**Estimate**	**Reference**	Δ	z**-score**
1	$4.7063(26) \times 10^{-10}$	4.706072e-10	2.542030e-14	0.096874
20	$4.46969192(47) \times 10^{-4}$	4.469691e-04	8.685338e-11	1.831710
40	$1.129572709(15) \times 10^{-2}$	1.129573e-02	1.679865e-09	11.024559
60	$7.435875869(22) \times 10^{-2}$	7.435875e-02	1.075293e-08	48.737161
80	$2.8840949115(43) \times 10^{-1}$	2.884094e-01	4.118825e-08	96.891117
100	$8.0968916849(86) \times 10^{-1}$	8.096891e-01	1.156258e-07	133.672155
120	$1.9174333291(11) \times 10^{0}$	1.917433e+00	2.727292e-07	244.117624
137	$3.5727881542(14) \times 10^{0}$	3.572788e+00	5.068474e-07	367.588934
140	$3.9550555805(15) \times 10^{0}$	3.955055e+00	5.607843e-07	380.453919
160	$7.4200401291(21) \times 10^{0}$	7.420039e+00	1.048215e-06	503.520894

Table 3.2: Energies of the Gaussian nuclear model as calculated with FEM and compared with *tabulated values* from [83]. All energies given in Hartrees.

Z	**Estimate**	**Reference**	Δ	z**-score**
1	$4.7075(24) \times 10^{-10}$	4.700000e-10	7.457202e-13	3.067855
20	$4.46890013(49) \times 10^{-4}$	4.468880e-04	2.013109e-09	40.904158
40	$1.129069158(44) \times 10^{-2}$	1.129069e-02	3.580371e-09	8.125451
60	$7.430213389(99) \times 10^{-2}$	7.430214e-02	-6.112159e-09	-6.187455
80	$2.880855998(14) \times 10^{-1}$	2.880856e-01	1.983463e-08	14.235654
100	$8.084723759(17) \times 10^{-1}$	8.084722e-01	1.458602e-07	86.928602
120	$1.9137452012(21) \times 10^{0}$	1.913745e+00	3.011955e-07	145.418463
137	$3.5645856426(30) \times 10^{0}$	3.564585e+00	4.825563e-07	161.807867
140	$3.9457096903(23) \times 10^{0}$	3.945709e+00	7.402734e-07	322.916540
160	$7.3990721532(26) \times 10^{0}$	7.399071e+00	1.053165e-06	408.224957

very good accuracy with the FEM-based approach and it indeed reproduces the exact energies if the basis is large enough. As the analytical approach only works for certain nuclear models (e.g. there are no known analytical solutions for the Gaussian nucleus) and require re-derivation for each new model, this demonstrates the power of the approach taken in this thesis, which is more generic and should work for any nuclear model, yet still yields machine-precision accuracy.

For completeness, Tables 3.1 and 3.2 tabulate the numerical values for the energies and how they compare to the tabulated reference values. Table 3.3 shows the comparison with the analytical energies for the hollow nucleus. It should also be noted that, as the analytic approach was re-implemented, this allows the calculation of reference values for other n and ℓ values, and so the table includes a comparison with the $n = 2$ values as well. As expected, the results are consistent between the analytic and FEM approaches.

Table 3.3: Energies of the hollow nuclear model as calculated with FEM and compared with directly calculated analytical energies for $n = 1$ (top table) and $n = 2$ (bottom table). All energies given in Hartrees.

Z	**Estimate**	**Reference**	Δ	z**-score**
1	$4.7063(26) \times 10^{-10}$	4.705996e-10	3.308465e-14	0.126082
20	$4.46969192(47) \times 10^{-4}$	4.469692e-04	1.904255e-11	0.401601
40	$1.129572709(15) \times 10^{-2}$	1.129573e-02	6.468781e-11	0.424531
60	$7.435875869(22) \times 10^{-2}$	7.435876e-02	7.480594e-11	0.339054
80	$2.8840949115(43) \times 10^{-1}$	2.884095e-01	3.228706e-11	0.075952
100	$8.0968916849(86) \times 10^{-1}$	8.096892e-01	8.640200e-11	0.099887
120	$1.9174333291(11) \times 10^{0}$	1.917433e+00	-1.282388e-10	-0.114785
137	$3.5727881542(14) \times 10^{0}$	3.572788e+00	-1.491571e-10	-0.108176
140	$3.9550555805(15) \times 10^{0}$	3.955056e+00	-4.583853e-10	-0.310983
160	$7.4200401291(21) \times 10^{0}$	7.420040e+00	-4.347385e-10	-0.208831

Z	**Estimate**	**Reference**	Δ	z**-score**
1	$5.9378(22) \times 10^{-11}$	5.882900e-11	5.493106e-13	24.730782
20	$5.58711891(77) \times 10^{-5}$	5.587119e-05	-4.476419e-13	-0.058366
40	$1.411972805(19) \times 10^{-3}$	1.411973e-03	9.350742e-12	0.499625
60	$9.294978097(39) \times 10^{-3}$	9.294978e-03	1.205080e-11	0.310169
80	$3.6052313945(60) \times 10^{-2}$	3.605231e-02	2.614797e-11	0.433689
100	$1.01216832650(87) \times 10^{-1}$	1.012168e-01	8.526513e-11	0.976146
120	$2.3970130852(16) \times 10^{-1}$	2.397013e-01	-2.660272e-11	-0.166148
137	$4.4665749415(19) \times 10^{-1}$	4.466575e-01	-1.364242e-11	-0.073067
140	$4.9445115194(24) \times 10^{-1}$	4.944512e-01	9.640644e-11	0.397063
160	$9.2769148153(27) \times 10^{-1}$	9.276915e-01	1.432454e-10	0.530625

3.4 Dirac equation with FEM

This section presents an accurate, reliable and extensible way of solving the radial Dirac equation using the finite element method, with a solver built on top of the HelFEM codebase [88]. While the focus is on the single-particle equation with an arbitrary nuclear potential, the method also generalizes to the many-particle mean-field case as well.

The aim is to determine, to machine precision, the eigenstates and eigenvalues of the radial Dirac equation (2.30)

$$\begin{pmatrix} mc^2 + V(r) & c\left(-\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right) \\ c\left(\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right) & -mc^2 + V(r) \end{pmatrix} \begin{pmatrix} P(r) \\ Q(r) \end{pmatrix} = E \begin{pmatrix} P(r) \\ Q(r) \end{pmatrix}$$

which in this and in the following sections will be in atomic units, where m is the mass of the particle in units of electron mass, distances are measured in Bohr radii a_0, energies in Hartrees, and $c = \alpha^{-1}$. The radial functions $P(r)$ and $Q(r)$ are related to the four-component orbitals according to Equation (2.28)

$$\psi(\boldsymbol{x}) = \frac{1}{r} \begin{pmatrix} P(r)\chi_{\kappa m}(\theta, \phi) \\ iQ(r)\chi_{-\kappa m}(\theta, \phi) \end{pmatrix}$$

which is a useful way of defining the orbitals because as the factor $1/r$ ensures that the solutions must obey the boundary conditions $P(0) = 0 = Q(0)$ at the origin. This can be imposed on the basis set level by excluding the left-most basis function that is non-zero at the origin.

As was mentioned in Section 3.2, the FEM basis also introduces another non-physical boundary condition at $r_{\max}$ where $P(r_{\max}) = 0 = Q(r_{\max})$. Like in the non-relativistic case, ideally $r_{\max} = \infty$, but this would either require some sort of special treatment or a different kind of basis set (such as Slater or Gaussian functions). However, again, this should not become an issue as the bound orbitals fall off exponentially far away from the origin, and therefore the eigenvalues and eigenstates should convergence well with increasing $r_{\max}$. But it is something that needs to be checked explicitly to make sure that the $r_{\max}$ is large enough for a given calculation.

FEM representation of the Dirac equation

As discussed in Section 3.2, the finite element method provides a set of functions on the domain $r \in [0, r_{\max}]$, to be used as a basis set for representing the orbitals. The solver for the Dirac equation was written in a way that it would be possible to use different basis sets for the components $P(r)$ and $Q(r)$.[8] As such, formally, there are two separate basis sets $\{b_i^{\mathrm{L}} \mid i = 1, \ldots, N_{\mathrm{L}}\}$ and $\{b_i^{\mathrm{S}} \mid i = 1, \ldots, N_{\mathrm{S}}\}$ for the large and small component $P(r)$ and $Q(r)$, respectively, which in general can have different numbers of basis functions (N_{L}, N_{S}). The two components are then represented as superpositions

$$P(r) = \sum_i^{N_{\mathrm{L}}} c_i^{\mathrm{L}} b_i^{\mathrm{L}}(r), \quad Q(r) = \sum_i^{N_{\mathrm{S}}} c_i^{\mathrm{S}} b_i^{\mathrm{S}}(r)$$

A general operator A acting on the two-component functions, consisting of parts that both act solely on the large and small component subspaces separately, but also of the parts that mix the

[8] While in the end it turned out to be unnecessary in this case, the ability to use different basis sets can be useful. For example, with a B-spline basis set, it appears to be necessary to use different orders of polynomials for the two components to ensure numerical stability [103, 102].

subspaces, can be represented as a $(N_\mathrm{L} + N_\mathrm{S}) \times (N_\mathrm{L} + N_\mathrm{S})$ block matrix. It can be arranged according to the large and small component subspaces as

$$A = \begin{pmatrix} A_\mathrm{LL} & A_\mathrm{SL} \\ A_\mathrm{LS} & A_\mathrm{SS} \end{pmatrix}$$

and the elements of the submatrices A_XY are calculated according to

$$(A_{XY})_{ij} = \int_0^{r_\mathrm{max}} \mathrm{d}r \ \left[b_i^X(r)\right]^* \ A_{XY} \ b_j^Y(r) \tag{3.5}$$

where $X, Y \in \{\mathrm{L}, \mathrm{S}\}$. The overlap matrix S is slightly special in that it is necessary to explicitly define that there is no overlap between the two subspaces, i.e.

$$S = \begin{pmatrix} S_{LL} & 0 \\ 0 & S_{SS} \end{pmatrix} \quad \text{with} \quad (S_{XX})_{ij} = \int_0^{r_\mathrm{max}} \mathrm{d}r \ \left[b_i^X(r)\right]^* b_j^X(r) \tag{3.6}$$

From Equation 2.30, the corresponding operator components of the radial Dirac Hamiltonian are

$$H_\mathrm{LL} = mc^2 + V(r), \quad H_\mathrm{LS} = c\left(-\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right)$$

$$H_\mathrm{SL} = c\left(\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right), \quad H_\mathrm{SS} = -mc^2 + V(r)$$

It is helpful to decompose these operators down further by defining a few generic operators

$$D_{XY} = \frac{\mathrm{d}}{\mathrm{d}r}, \quad R_{XY}^{(n)} = r^n, \quad \Phi_{XY} = \frac{V(r)}{Z\alpha} \tag{3.7}$$

acting on any combination of subspaces, with the matrix elements calculated according to (3.5). Φ_{XY} represents the normalized nuclear potential $\Phi(r)$ corresponding to $V(r)$ defined by (2.31), although in atomic units the relationship becomes

$$V(r) = Z\alpha\Phi(r)$$

By fully shifting into atomic units by replacing the speed of light constant, the matrix for the Hamiltonian in the FEM basis can be constructed as

$$H = \begin{pmatrix} m\alpha^{-2}S_\mathrm{LL} + (Z\alpha)\Phi_\mathrm{LL} & -\alpha^{-1}D_\mathrm{SL} + \kappa\alpha^{-1}R_\mathrm{SL}^{(-1)} \\ \alpha^{-1}D_\mathrm{LS} + \kappa\alpha^{-1}R_\mathrm{LS}^{(-1)} & -m\alpha^{-2}S_\mathrm{SS} + (Z\alpha)\Phi_\mathrm{SS} \end{pmatrix} \tag{3.8}$$

This form of the Dirac Hamiltonian is useful because the matrix elements of the component operators (3.7) can be calculated in advance, independently of the various physical properties of the system (such as the nuclear charge). Computing the eigenvalues for a whole range of parameters, say when scanning in Z, is now just a matter of summing up the constituent matrices with the correct coefficients. The only exception for that is the matrix for the normalized potential Φ_{XY}, which may have to be recomputed often as the shape of the nuclear charge distribution and therefore the potential varies with the nuclear charge.

At this point, the Hamiltonian matrix (3.8) and the overlap matrix (3.6), could be passed to a generalized eigenvalue solver to determine the eigenvalues and eigenvectors. However, the implementation again uses Löwdin orthogonalization [108] instead. For that, we need to calculate

the half-inverse of the overlap matrix which, since the overlap matrix is block-diagonal, reduces to the half-inverses of the component overlap matrices

$$S^{-1/2} = \begin{pmatrix} S_{\mathrm{LL}}^{-1/2} & 0 \\ 0 & S_{\mathrm{SS}}^{-1/2} \end{pmatrix}$$

As was mentioned in Section 3.2, this effectively transforms the basis into an orthonormal basis and the transformed Hamiltonian matrix

$$\tilde{H} = S^{1/2} H S^{-1/2}$$

can be diagonalized normally without having to worry about a generalized eigenvalue problem. The eigenvectors ψ in the original basis can be recovered from the eigenvectors in the orthonormalized basis $\tilde{\psi}$ via

$$\psi = S^{-1/2}\tilde{\psi}$$

which is necessary for graphing the orbitals, for example.

Choice of grid

Before embarking on the journey of trying to get relativistic energies and orbitals to converge, it is important to determine the most optimal grid for solving the relativistic one-particle problem. In general, relativistic orbitals are more concentrated near the origin relative to their non-relativistic counterparts. Therefore the exponential grid, which allocates elements closest to the origin, should still be the most suitable one. However, it might be beneficial to distribute the element boundaries closer to the origin because of the relativistic shift of the orbitals.

Like when benchmarking the method for the non-relativistic finite nucleus case, the energy shift when doubling the number of elements will be used as the convergence measure. As before, the goal is to find a region where the energy appears to be most converged with the fewest number of elements.

Point nucleus

Figure 3.16 shows the convergence measure for the point nucleus, with smaller values indicating faster convergence. At very low nuclear charge, low values of the z parameter achieve good convergence. However, already in the $Z = 30-50$ range, it becomes clear that the non-relativistic recommendation of $z = 2$ is no longer suitable, and this issue becomes even more severe at higher nuclear charge. Based on these figures, the recommendation when solving the Dirac equation for a point nucleus would be to use an exponential grid with $z = 6$. This should guarantee a similar convergence behaviour for all Z values, including close to the critical nuclear charge $Z = \alpha^{-1} \approx 137$.

It should also be noted that achieving convergence becomes very difficult close to the critical nuclear charge, as can be seen from the $Z = 137$ graph — even with 160 elements it is nowhere near to being converged there. This is expected, as the orbitals approach the $\mathrm{e}^{-\lambda r}$ shape in the limit of $Z \to \alpha^{-1} \approx 137$, leading to a very sharp cusp near the origin. No basis set that imposes the $P(r) = 0 = Q(r)$ boundary condition at the origin fully describes that behaviour. However, it can be concluded that the results for $Z \leq 120$ are good and reliable. It would also be possible to extend the range higher into Z with additional elements, but it would be necessary to carefully make sure that the results actually converge.

In principle, it is technically possible to perform the numerical calculations for $Z \geq \alpha^{-1}$ and it would lead to some energies and orbitals. However, such calculations will not be particularly

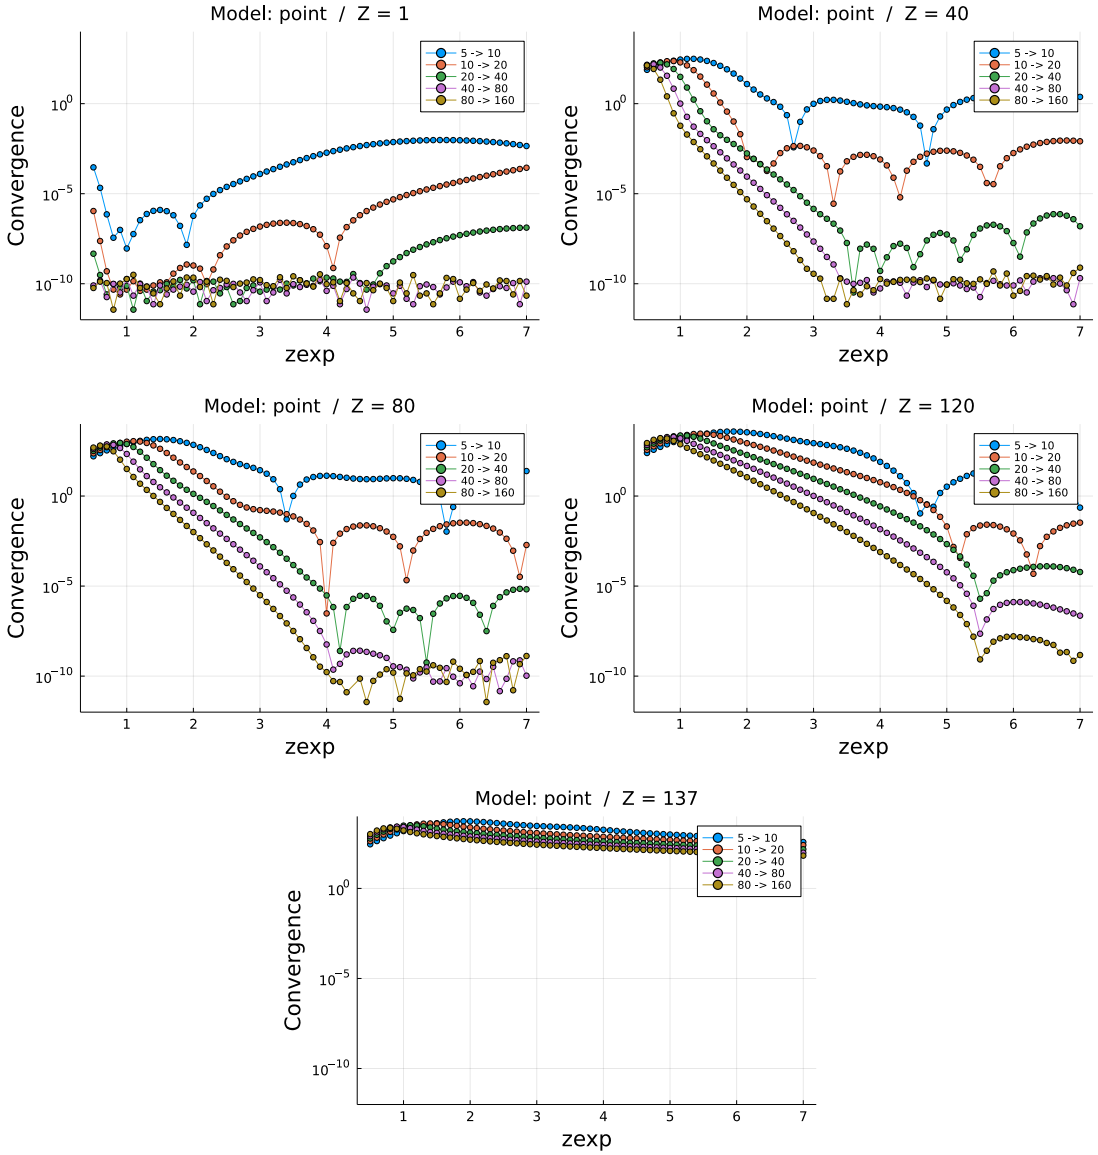

Figure 3.16: Convergence of the $1s$ energy of the Dirac equation for the *point nucleus* for various Z values. The calculations were performed with 6th order polynomials and $r_{\max} = 40\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

meaningful from a physical or mathematical point of view, as the singularity in the $1/r$ operator, which causes the Dirac equation to break down at $Z = \alpha^{-1}$, can not be exactly represented by a finite basis of the type we are using. This is evident from the convergence problems that become more and more severe when approaching $Z \to \alpha^{-1} \approx 137$ from below.

Finite nuclei

The grid benchmarks for finite nuclear models are less dramatic, with the scans of the z parameter showing in Figure 3.17 for the *hollow nuclear model*. The initial behaviour at lower nuclear charges is similar to the point nucleus, except that it does not seem necessary to make z as high as for the point nucleus — the convergence rate appears to more or less flatten out at around $z = 3-4$. However, unlike in the point nucleus case, it is also possible to converge high Z values, such as $Z = 137$. This, of course, is not unexpected, as the finite nuclear models do not have the problematic singular nature of the $1/r$ potential.

Additional convergence graphs (for additional nuclear models and Z values) are presented in Appendix C, and the same insights to hold for all finite nuclear models. With that in mind, the recommendation is to use $z = 3$ as the grid when performing calculations with physical finite nuclear models.

Energies

Having determined a reliable grid for the relativistic problem, the next step is to determine if it is possible to consistently converge the energies. The algorithm for determining the converged energies will be the same as in the non-relativistic finite nucleus case described in Section 3.3. In a nutshell, the energies for all element values are calculated and then a sliding window of a fixed width (20 points, like in Section 3.3) is run over the number of elements, and the range where the variance of the energy is the smallest is picked.

Figure 3.18 illustrates how the energy converges with an increasing number of elements by plotting the change in energy as the number of elements is increased. The energy quite systematically converges to floating-point precision, which in this case is approximately 10^{-10} Ha. This is consistent with the use of 64-bit floating-point numbers where the approximately $10^{-14} - 10^{-15}$ relative floating-point precision if taken relative to $mc^2 \approx 1.8 \times 10^4$ — the energy scale of the Dirac problem in atomic units — comes out to roughly that order of magnitude. The seemingly discontinuous jumps in the energy convergence are caused by the element boundaries crossing the extra boundary added for discontinuous nuclear potentials (to accurately represent the discontinuity; see Section 3.3).

The energies determined by the algorithm are shown in Figures 3.19, 3.20, and 3.21. In each case, the order of the polynomials and the z parameter of the exponential grid is varied to show that the method reliably achieves consistent converged results independent of the basis parameters for both finite and point nuclei. The graphs for all cases are available in Appendix C.

Figure 3.20 confirms that for the point nucleus and near the critical charge $Z = \alpha^{-1} \approx 137$, it is difficult to achieve convergence. The determined uncertainty in this case is on the order of 1 Ha, whereas for converged cases it is on the order of 10^{-10} Ha. On the other hand, Figure 3.21 shows that the critical charge is not a problem in the finite nuclear case. But picking a z parameter that is too high, which would concentrate the elements very close to the origin, is not a good idea for the finite nuclear case.

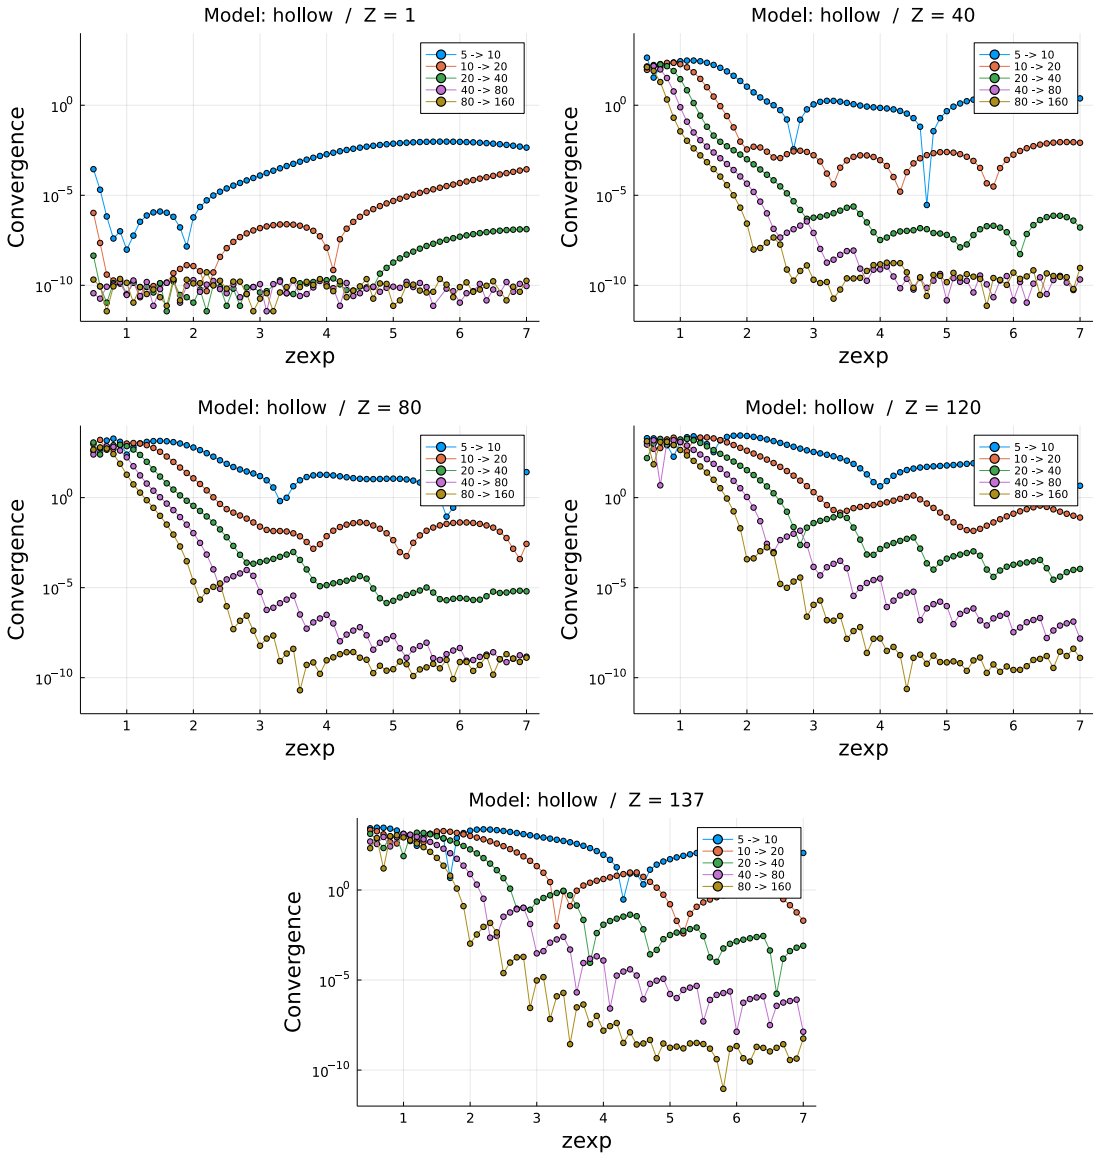

Figure 3.17: Same as Figure 3.16, but for the *hollow nuclear model*.

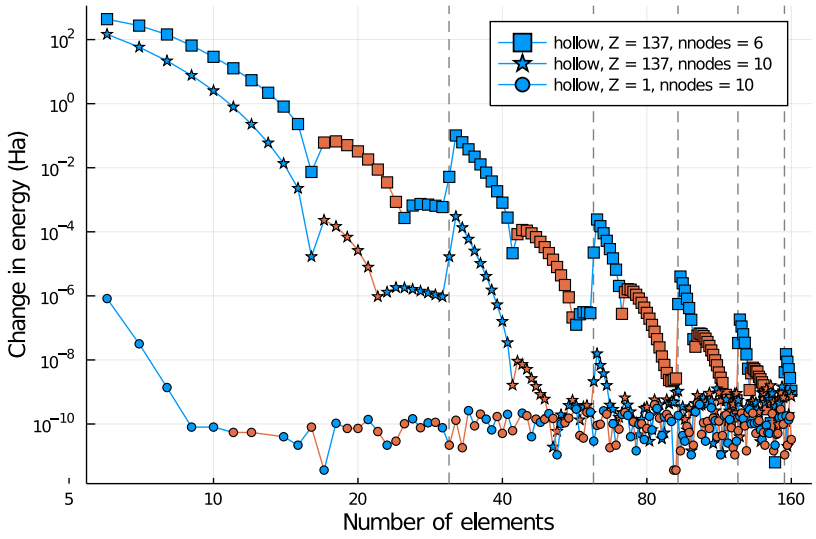

Figure 3.18: Examples of the convergence behaviour of the 1s energy for finite nuclei for different Z values and basis polynomial orders as the number of basis elements increases. The different colours indicate whether the change in energy is positive or negative (red or blue, respectively). The vertical grey dashed lines show where an element boundary from the exponential grid crosses over the manually inserted extra boundary as the element boundaries get pushed towards the origin as additional elements are added.

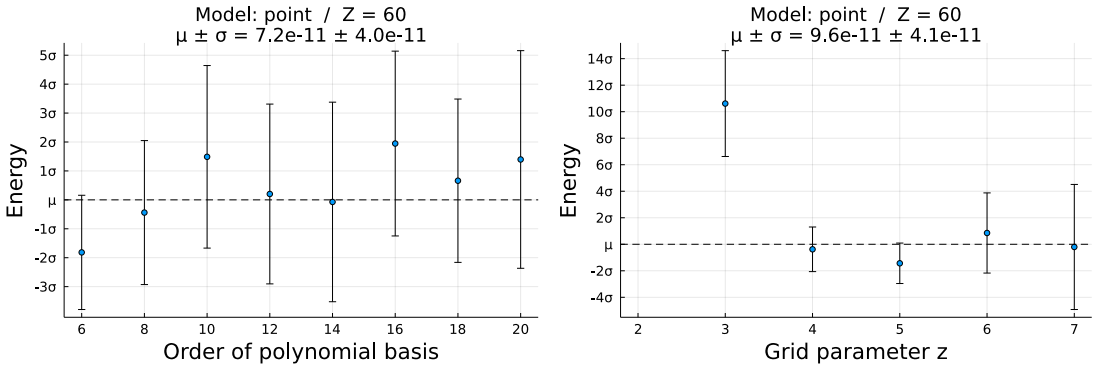

Figure 3.19: Consistency of energy values for the point nuclear model for $Z = 60$. Left column is consistency between runs with different basis polynomial orders and right columns is varying the x parameter of the exponential grid, leading to a different grid of element boundaries.

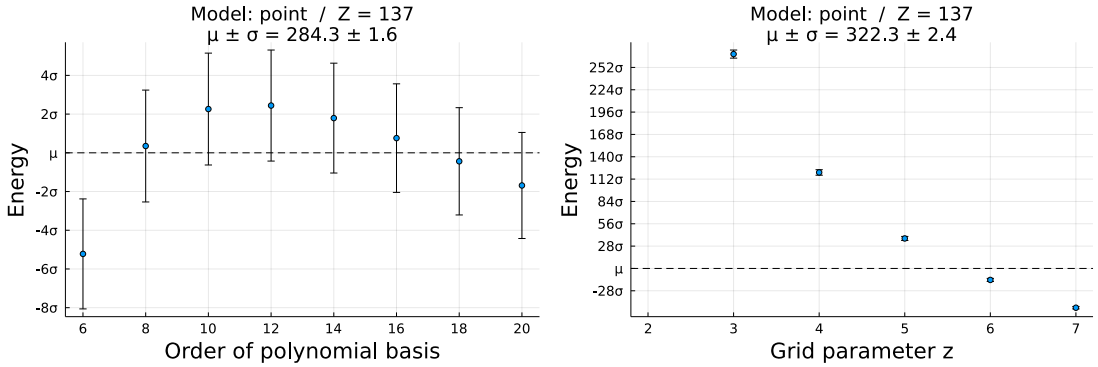

Figure 3.20: Same as Figure 3.19, but for $Z = 137$, showing the failure of convergence for the point nucleus near the critical charge (note the extremely high uncertainty on the scale of 1 Ha).

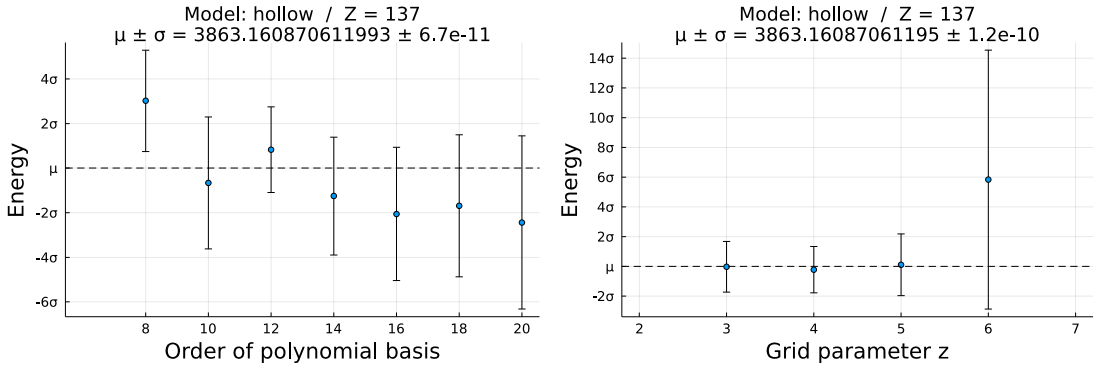

Figure 3.21: Same as Figure 3.19, but for the hollow nuclear model at $Z = 137$.

3.5 Diving into the negative energy continuum

While having a way to calculate the eigenenergies and eigenstates of the Dirac equation to high accuracy is excellent, the true test of the finite element method is to see whether it can be used to explore the various edge cases of the Dirac equation as well. The first application of the method is trying to describe what happens to the bound-state solutions once the energy of the 1s state drops below $-mc^2$ and becomes degenerate with the negative energy states (or "dives into the sea of negative energy continuum states").

It should be pointed out that the discussion of "diving" only makes sense in the context of a finite nucleus, as for a point nucleus (i.e. a bare $1/r$ potential) the Dirac equation breaks down at $Z \geq \alpha^{-1} \approx 137$, where the energy becomes zero is nowhere near the negative energy continuum yet. In the case of a finite nucleus, however, there is no singularity at the origin and the equation can be solved without any problems even as the energy falls below zero. But the correct treatment of the state when it becomes degenerate with the continuum is unclear, and the goal here is to see if it is possible to recover a reasonable bound state from the continuum in the numerical FEM approach.

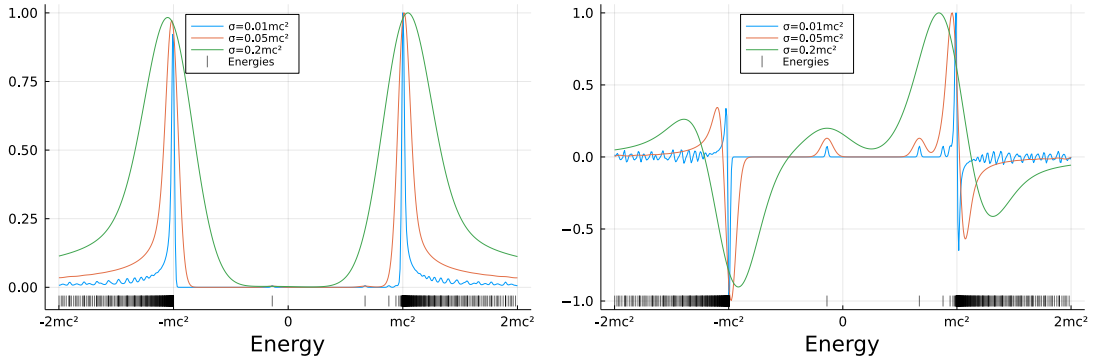

Figure 3.22: Graphs showing the density of states for $s_{1/2}$ states as a function of energy, for different Gaussian widths (using the $Z = 150$ case as an example). The graph on the left shows the absolute density, whereas the graphs on the right shows the difference relative to the free particle case ($Z = 0$; i.e. illustrating how the density has shifted due to the introduction of an attractive central potential). The ticks below either graphs indicate where the discrete energies that we are smoothing are located. The y-axes have been scaled to make the different cases visually comparable.

Gaussian smoothing of the density of states

As the numerical calculations are always performed with a finite basis set, and also formally the radial coordinate has been cut off, imposing a confinement-like boundary condition at some finite radius, the spectra obtained from the finite element calculations will always be made up of discrete points and do not exhibit any continua. However, the method is able to recover an approximation to the positive and negative energy continua (above mc^2 and below $-mc^2$, respectively), where the density of states is much greater than in the intermediate bound state region at $-mc^2 < E < mc^2$. This is visible on the bottom parts of Figure 3.22, where each vertical tick represents one discrete eigenstate of the discretized Hamiltonian.

As visually interpreting the discrete points is difficult, it would be useful to have a way to visualize the changes in the distribution of states, for both the continua and the true discrete bound states. An option for that is to smudge the discrete states by replacing each of them, which all can formally be thought of as a delta distribution, with a narrow Gaussian distribution that has a finite width. In other words, given a set of discrete eigenenergies $\{E_i\}$, the corresponding Gaussian-smoothed probability density $g(E)$ is defined to be

$$g(E) = \sum_i \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2}\left(\frac{E - E_i}{\sigma}\right)^2\right]$$

where σ is the chosen width (in units of energy) of each of the Gaussian distributions applied to each point. Technically, the value of this function measures the number of states per unit energy and has a total integral equal to the number of discrete data points (which in turn is equal to the number of basis functions).

Figure 3.22, plotting the smoothed densities for various widths, illustrates what Gaussian smoothing achieves. It provides a useful way to visualize the density of the states, including the contrast between the bound states and continuum states. What can also be seen on those figures is that having a width σ that is too small (i.e. Gaussians that are too narrow) will lead to wiggles in the continuum, as the width will be smaller than the distance between the discrete

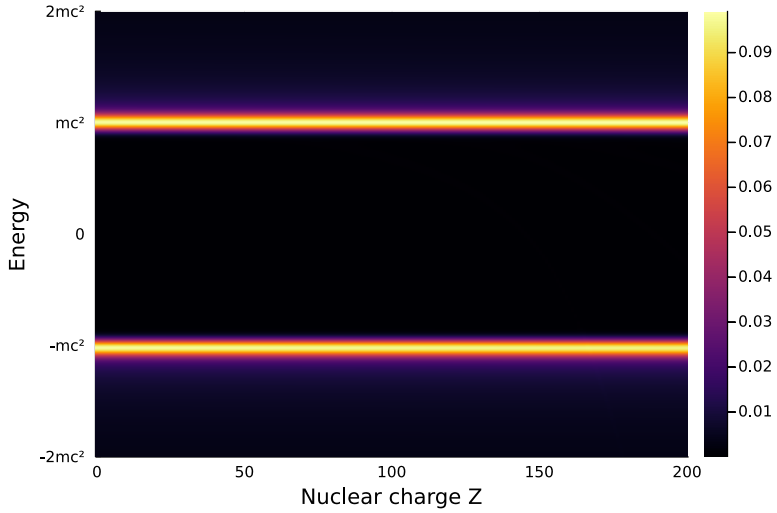

Figure 3.23: A heatmap of the Gaussian-smoothed density of states from the numerical FEM calculation as a function of Z. There is no visible change with respect to the nuclear charge since the changes in the density are, relatively speaking, very small.

points. On the other hand, Gaussians that are too wide will smooth out all the detail, especially the bound states.

What Figure 3.22 also shows is that the absolute density is not particularly illuminating, as what is mostly visible there is a large density near $\pm mc^2$, that rapidly drops off. This particular distribution of the states is a consequence of solving a particle-in-a-box problem and using a finite basis set, with less oscillatory continuum states with lower (absolute) energy being easier to represent with the basis set. Most importantly, however, the bound states are essentially invisible, although very tiny bumps can be seen between the two continua if one looks very closely.

However, the density of the continua should not change that dramatically when the central potential changes, and this is demonstrated in Figure 3.23. Therefore, rather than plotting the absolute value of the density, it turns out to be advantageous to look at the change in the density of states relative to the pure-continuum case ($Z = 0$) instead. That is, if $g(E; Z)$ is the density for the nuclear charge Z, then it is more useful to plot

$$\delta g(E; Z) = g(E; Z) - g(E; 0)$$

instead, shown on the right in Figure 3.22. What is immediately clear is that the bound states show up much more clearly. There are also some shifts in the continuum: the upper continuum is slightly pulled downwards, which is intuitive as the bound states mostly consist of positive-energy continuum states. However, the negative continuum is also slightly affected, with the density slightly reduced overall and pushed downwards, deeper into the negative energies. Most of the states that appear to affect the bound states are the low energy continuum states (i.e. near $\pm mc^2$).

We can use these insights to visualize how the density of states changes with the nuclear charge. First, Figure 3.24 shows the density of states differences for a few different Z values. What shows up very clearly are the bound states, moving down in energy with the increasing Z, hitting the negative continuum, and then continuing towards negative infinity. This is even more

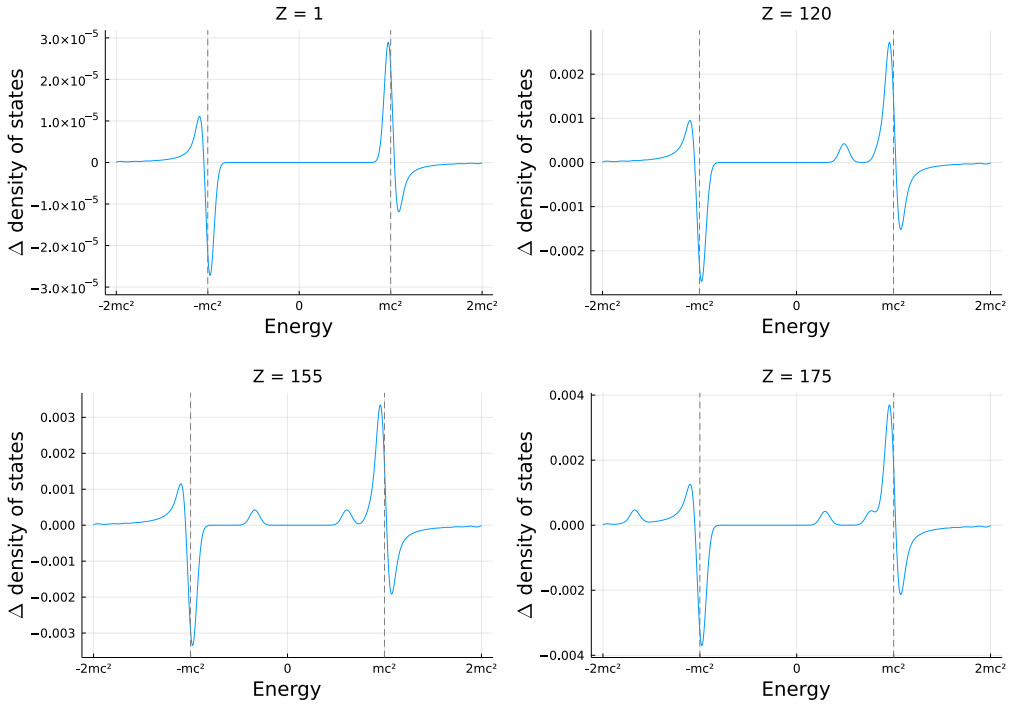

Figure 3.24: Changes in the density of states relative to the pure continuum case ($Z = 0$) for various Z values.

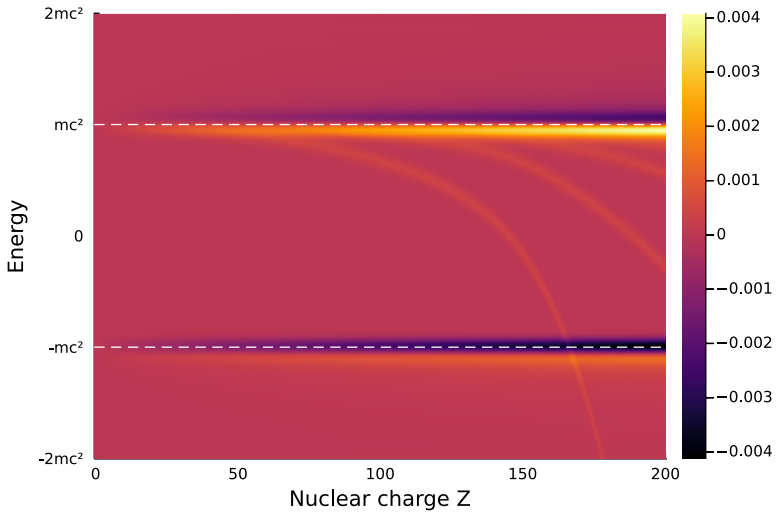

Figure 3.25: A heatmap of the change in the Gaussian-smoothed density of states relative to the pure continuum ($Z = 0$) case for a numerical $s_{1/2}$ calculation. The dark areas are where the density is reduced, the red is where it is unchanged, and the yellow parts are where the density has increased. The diving of the $1s_{1/2}$, $2s_{1/2}$, and $3s_{1/2}$ can clearly be seen as distinct lighter-colored curves.

clearly visible in Figure 3.25, where the spectra have been calculated for a dense grid of Z values. These numerical results suggest that after the $1s_{1/2}$ state has dived into the negative continuum, it continues moving downward and is a distinct additional state on top of a continuum that has not been disturbed that much.

Tracing bound states in the negative energy continuum

Technically, the calculations with a finite basis set simply yield a list of eigenenergies and eigenvectors. As such, while it visually appears that it would be easy trace a bound state in the spectrum of the Dirac operator, including after it has dived into the negative energy continuum, it is actually anything but. It happens that, for lower Z values before the diving, the number states in the negative continuum is equal to N_S — the number of basis functions used to describe the small component. This means that the few lowest bound states can very easily be picked out from the spectrum by simply discarding the N_S lowest energies in our spectrum. However, this obviously is not the case anymore once the energy of the $1s_{1/2}$ state goes below $-mc^2$ and the bound state becomes degenerate with the states in the negative energy continuum. At that point, there is no immediately obvious way to distinguish the bound state from the other continuum states, and it is necessary to figure out a strategy for identifying the bound states once they are located in the continuum.

What follows stems from the hypothesis that, the character of the bound state should remain consistent and change smoothly with the increasing Z value, both qualitatively and quantitatively. If this is true, vector overlaps between the states from calculations with different Z values can be used to determine the correspondences between the states. To be more quantitative, let $\{|k; Z\rangle\}$ denote the set of all the eigenstates of a Hamiltonian with nuclear charge Z in some basis. k is an arbitrary index for labelling the different eigenstates for a particular nuclear charge. Further, let k_0 be the index of the $1s_{1/2}$ state for some nuclear charge Z_0 in the pre-diving region (i.e. $Z_0 < Z_{\mathrm{cr}}$, where Z_{cr} is the critical Z at which diving occurs; as discussed above, in this region we can easily identify the index of the state). The strategy for finding the state in the continuum that corresponds to the pre-diving state, therefore, would involve looking at the probabilities

$$\left|\langle k; Z|k_0; Z_0\rangle\right|^2 \tag{3.9}$$

However, finding a reliable strategy turned out to be trickier than initially thought, and for that reason it is worth discussing a few of the unsuccessful strategies as well, before describing the one that actually worked.

Step-by-step strategy

The immediate concern with a naive strategy of using the maximum overlap probability (3.9) directly is that the character of the $1s_{1/2}$ state changes significantly as Z increases, and so the overlap with a reference state far away in Z is no longer helpful. This issue is even more severe due to the diving, as the reference state will always be from the pre-diving region, and one would expect that, in the post-diving region, the (near) degeneracies with the negative energy continuum states would considerably affect the character of the bound state.

The simplest way to work around that issue was to try to simply step, one point at a time, from the reference state on a reasonably fine grid in Z. At each point Z_n, the state with the maximum overlap probability with the previously determined bound state at Z_{n-1} is declared to be the bound state. However, as can be seen in Figure 3.26, this does not work particularly well. The trace, once it gets deep enough into the negative continuum, starts following one of the continuum states instead. The reason for that can be seen in on the right panel of the same

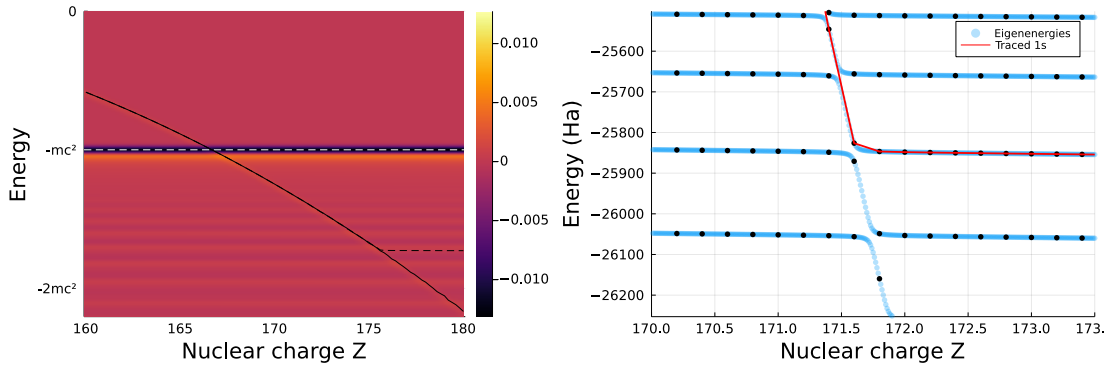

Figure 3.26: Figures showing the tracing of the $1s_{1/2}$ state by using the overlap with the state found for the previous Z value (or averaged overlap of previous five states for the dashed line) on a fine Z grid. Once the state dives deep enough, the trace fails to follow the correct path due to the trace following an avoided level crossing (shown on the right, where the energy axis has been zoomed in close to where the tracing fails; the points show the exact eigenenergies; black points are the ones used in the tracing, but the blue lines were calculated to get a finer grid between the points to better illustrate the avoided crossing).

figure: there are avoided crossings between the negative continuum states and the diving $1s_{1/2}$ state, likely due to the finite r cutoff. Of course, in the case of an avoided crossing, the state will change character from the $1s_{1/2}$ bound state into a continuum state but will do so continuously in terms of the Z value. This means, perhaps counter-intuitively, that making the grid in Z finer when trying to trace the state will actually make it harder to follow the true $1s_{1/2}$ state.

One could hope that perhaps averaging over the states from several previous points in Z would help. It usually does, a little, as illustrated by the dashed line in Figure 3.26, but this does not scale: as the bound state falls deeper and deeper into the continuum, the avoided crossings get more and more severe, and so the strategy will fail eventually, independent of how many previous points are included in the algorithm.

Single-reference strategy

While, as was just said, the most naive strategy of using a single reference state at a fixed Z value would likely not work, it is interesting to see how it fails. As can be seen in Figure 3.27, this approach actually usually fails before even entering the diving regime. The reason for the failure can be seen in Figure 3.28, which shows the overlap probabilities that were used to determine the $1s_{1/2}$ state at another Z value. It confirms the suspicion that the character of the $1s_{1/2}$ state, as Z increases, changes so significantly that it is no longer comparable to the reference state. While for $Z = 100$ the $1s_{1/2}$ state (dot with the lowest energy between the two continua) has the highest probability and would therefore be picked up by the algorithm, that is no longer the case for $Z = 140$. Once Z becomes large enough, the reference $1s_{1/2}$ will actually have a greater overlap with the $2s_{1/2}$ state than with the $1s_{1/2}$ of the lower Z value, and so the trace would jump to the $n = 2$ line. This is not wholly unexpected — as Z increases, the bound states contract significantly, making the higher n states look more like the reference state in terms of the spacial distribution of the orbital.

The overlaps on the right in Figure 3.28, specifically the $Z = 170$ case where the $1s_{1/2}$ has dived into the continuum, show another interesting feature. As the bound state dives and be-

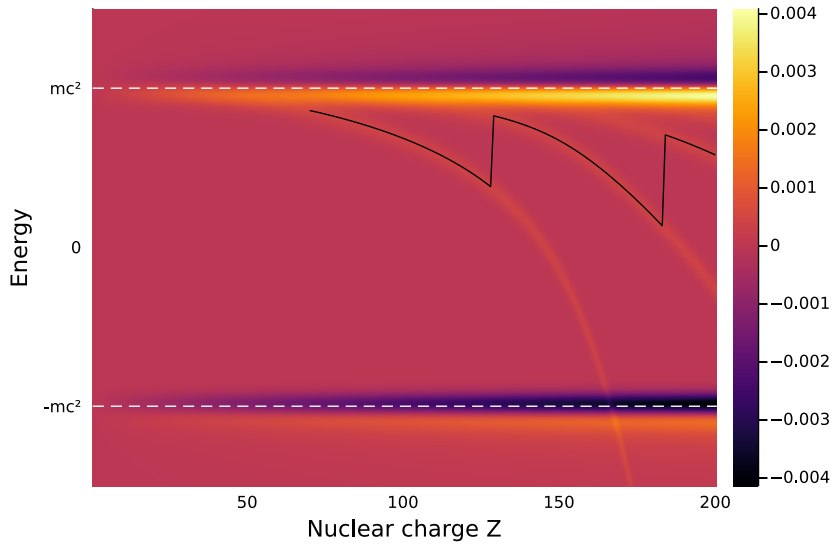

Figure 3.27: An example of the failure of the single pre-critical reference tracing strategy — the trace starts at $Z = 70$, and initially follows the $1s_{1/2}$, but then jumps to follow the $2s_{1/2}$ and eventually the $3s_{1/2}$ state.

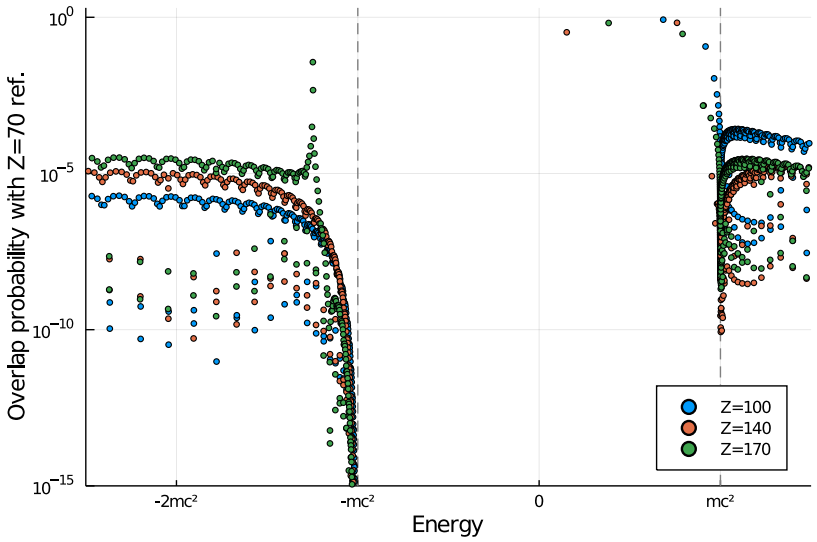

Figure 3.28: Related to Figure 3.27, showing the the overlap probability $|\langle k_0, Z_0|k, Z\rangle|^2$ between the reference $1s_{1/2}$ state for $Z = 70$ and all the states in the given energy range for other Z values.

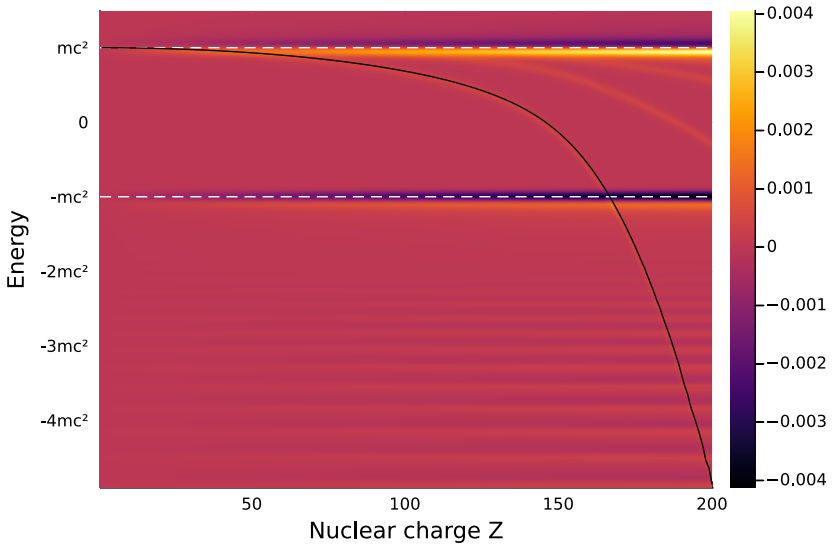

Figure 3.29: A successful tracing of the $1s_{1/2}$ state into the negative continuum. The tracing was done by looking at the overlap with the $1s_{1/2}$ for $Z = 1$, but by always discarding all the eigenstates with energy greater than the last found energy.

comes degenerate with the negative continuum states (at least, approximately, since numerically the states are still discrete), it appears to spread across multiple continuum states. From the perspective of the state tracing algorithm, this is quite problematic since it reduces the absolute value of the overlap by spreading it across multiple states. It is, therefore, very likely that simply trying to search for the maximal overlap value will lead to spurious matches with undived bound states and other states.

Single-reference strategy with energy filtering

While a strategy that uses a pre-diving reference state does not work well at all, it can be observed that when it does fail, it is because it starts spuriously matching a bound state with higher energy (e.g. $2s_{1/2}$ or higher). That, together with the fact that the energy of the bound state gets monotonically lower and lower with increasing Z, leads to an idea for a slight modification of the algorithm: what if all the states with an energy higher than the last found are just discarded immediately?

Figure 3.29 shows the result. As the algorithm is not allowed to jump into the $2s_{1/2}$ and higher states, it manages to perfectly trace the diving of the $1s_{1/2}$ state. While the graph only shows the diving down to $-5mc^2$ and $Z = 200$, the algorithm works well also beyond that point, and its performance does not appear to be affected by the choice of the basis.

It should be noted, however, that it does not work that well for states with higher n quantum number (e.g. tracing $2s_{1/2}$). In those cases, the trace will eventually jump down to the lower energy bound states (i.e. the trace for $2s_{1/2}$ will start following the $1s_{1/2}$ instead). One potential way around that might be to assign a weight to the energy difference from the previous match so that large jumps away from the already found trace would be penalized.

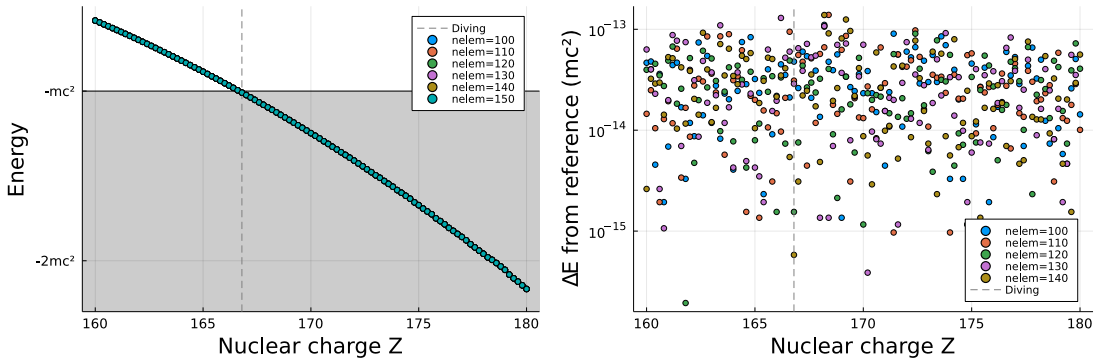

Figure 3.30: Convergence of the trace of the energy of the $1s_{1/2}$ state near diving, as we systematically increase the $r_{\max}$ value of the basis set by adding in extra elements at the end on the same exponential grid. The graph on the left shows the absolute energy for various numbers of elements, where as the graph on the right shows the energy difference from the case with the highest number of elements (and therefore, correspondingly, the largest $r_{\max}$ value).

Convergence of the traced energy

Now that it is possible to trace the $1s_{1/2}$ state into the negative energy continuum, it is necessary to make sure that this is actually physically meaningful by checking if the trace is stable and converges systematically with increasing basis and box sizes.

Figure 3.30 shows the trace for various $r_{\max}$ values, which impose the boundary condition at large but not infinite r. To check this systematically, rather than just stretching the same basis out over a larger range of r values, the domain was extended by adding new elements to the end of the basis, essentially by evaluating the exponential grid Eq. (3.4) at $i > N$. The original element boundaries (and also the primitive polynomial basis) are kept constant, so the only thing we gain is the ability to represent states farther out on the r-axis. What can be seen from the figure, especially the one on the right, is that the trace essentially does not change with increasing box size, with the difference attributable to numerical noise.

Figure 3.31, where the box is kept the same, but the basis is improved by making the grid denser, paints a slightly different picture, however. In the pre-diving region, while with 50 elements the energies have not quite converged, with more elements they do converge to machine precision. However, in the post-diving region, we get a relatively large noise, on the order of $10^{-3}\ mc^2$. However, it does look like this is noise and not a lack of convergence as such, as the energy fluctuates around some central value. It appears to be caused by the level crossings with the (numerical) continuum states and due to the smearing of the state across multiple continuum states. The energy differences between the continuum states are consistent with the magnitude of the numerical noise.

Orbitals in the continuum

Finally, in a numerical approach, it is extremely easy to reconstruct the P and the Q components of the orbitals as functions on r. Figure 3.32 shows two example cases of the $1s_{1/2}$ state, one before the diving and one after, and illustrates the general behaviour of the orbitals. In the pre-diving region the orbital is a clean bound state, and if one would look at how it changes as a function of Z, it starts out on the right and then smoothly contracts towards the left end of

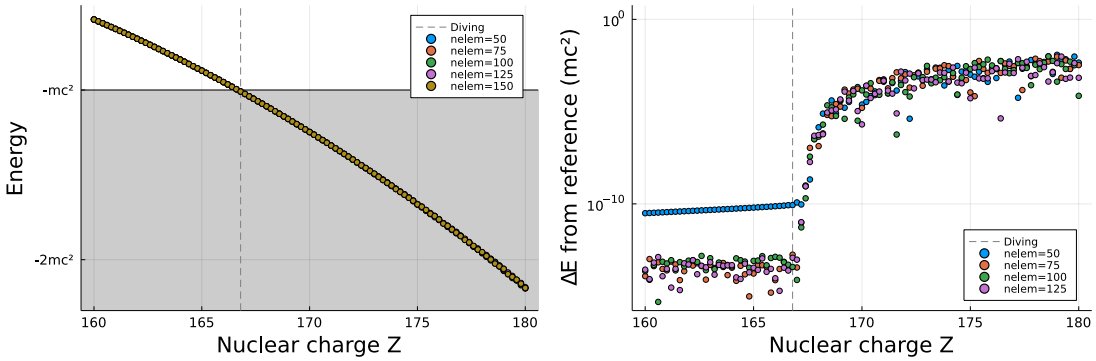

Figure 3.31: Convergence of the trace of the energy of the $1s_{1/2}$ state near diving, as we increase the basis size by adding elements, but keeping the $r_{\max}$ value constant (i.e. improving the representation of the orbitals). The graph on the left shows the absolute energy for various numbers of elements, where as the graph on the right shows the energy difference from the case with the highest number of elements.

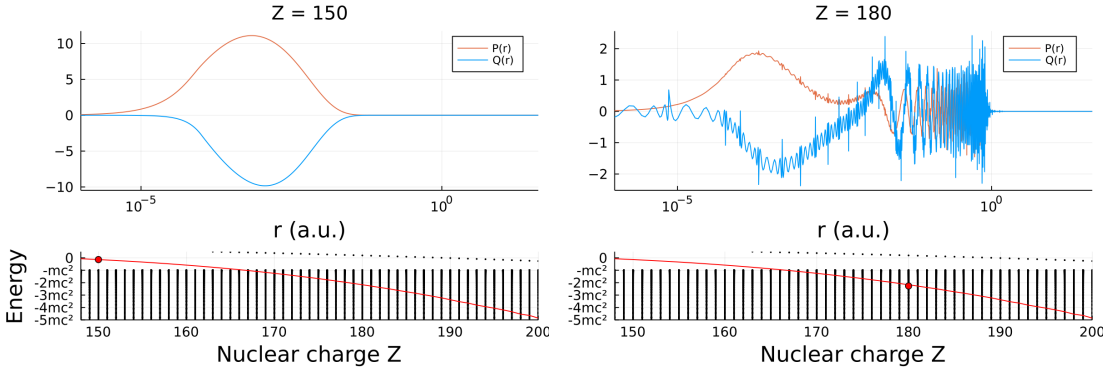

Figure 3.32: Examples of wavefunctions of the $1s_{1/2}$ eigenstates as calculated with FEM, in pre-critical ($Z = 150$; on the left) and post-critical ($Z = 180$; on the right) nuclear charge regions. Pre-critical orbitals are pure bound states with clean graphs, whereas after diving into the negative continuum the states start mixing in continuum oscillations. The graph on the bottom indicates which state on the trace is being plotted. For clarity, the right-most flat end of the right figure is also part of the numerical solution, just with a tiny amplitude.

the graph (towards the origin).

However, once in the continuum, the bound state starts mixing in the continuum solutions. This is expected, since the calculations are purely numerical, and there is no attempt to impose any particular character on the states, such as exponential decay. So degenerate states end up being mixed with rather arbitrary weights. It is also consistent with the avoided crossings observed earlier in the energies, which are likely due to numerics rather than physics. These crossings exacerbate the problem since it means that we get mixing not just when a true degeneracy occurs but even when the states have slightly different energies. All in all, though, an important observation here is that, while the state contains oscillations, it also retains a clear bound state component to it, and the energy as a function of Z remains relatively smooth.

3.6 Energy-projected Dirac equation

This section shows how the finite element method can be applied to a to a modified version of the single-particle Dirac Hamiltonian: the so-called *no-pair external field Dirac Hamiltonian*

$$h_+ = \Lambda_+(h_\mathrm{D} + V_\mathrm{ext})\Lambda_+ \tag{3.10}$$

where Λ_+ is the projection operator to the free-particle positive energy subspace of the free-particle Dirac Hamiltonian h_D. The motivation for such an equation is simple: a way to avoid the problem of the bound states diving into the negative energy continuum at high Z values (when solving the Dirac equation with a finite nuclear model) is to just remove the offending negative continuum. Effectively the projection operator restricts the Hilbert space to the subspace defined by the positive energy continuum states, leading to different eigenvalues and eigenstates compared to the original Hamiltonian.

As long as the Hamiltonian h does not have any zero eigenvalues, a projection operator to positive or negative energy subspaces can formally be written as

$$\Lambda_\pm = \frac{1}{2}\left[1 \pm \frac{h}{|h|}\right]$$

The eigenvalues of the free-particle Dirac Hamiltonian h_D are $|E| \geq mc^2$ [11], so it is possible to apply it in this case. However, evaluating the projection operator directly in real space is tricky. Instead, it is more convenient to write the operator in momentum space, where the free-particle Dirac Hamiltonian becomes a $\boldsymbol{p}$-dependent matrix multiplication operator $\hat{h}_\mathrm{D}(\boldsymbol{p}) = \alpha \cdot \boldsymbol{p} + m\beta$. As $|\hat{h}_\mathrm{D}(\boldsymbol{p})| = \sqrt{m^2c^4 + \boldsymbol{p}^2}$, the projection operator in momentum space can be written as

$$\Lambda_\pm(\boldsymbol{p}) = \frac{1}{2}\left[1 \pm \frac{c\boldsymbol{\alpha} \cdot \boldsymbol{p} + \beta mc^2}{E(\boldsymbol{p})}\right]$$

where $E(\boldsymbol{p}) = +\sqrt{m^2c^4 + \boldsymbol{p}^2}$. [109] Since the operators in real and momentum space are related by a Fourier transform $\Lambda_\pm = \mathcal{F}^{-1}\Lambda_\pm(\boldsymbol{p})\mathcal{F}$, it is now possible to represent the projection operator with the real-space operator kernel

$$\Lambda_\pm(\boldsymbol{x}, \boldsymbol{y}) = \frac{1}{(2\pi)^{\frac{3}{2}}}\mathcal{F}^{-1}[\Lambda_\pm(\boldsymbol{p})](\boldsymbol{x} - \boldsymbol{y}) = \frac{1}{2}\frac{1}{(2\pi)^3}\int\left[1 \pm \frac{c\boldsymbol{\alpha} \cdot \boldsymbol{p} + \beta mc^2}{E(\boldsymbol{p})}\right] \mathrm{e}^{i\boldsymbol{p}\cdot(\boldsymbol{x}-\boldsymbol{y})}\,\mathrm{d}^3\boldsymbol{p}$$

which could be integrated to formally obtain the full real-space operator.

This projected Dirac Hamiltonian (3.10) can be traced back all the way to Bethe and Salpeter [109], and is therefore sometimes referred to as the Bethe-Salpeter operator [110]. However, the more illustrative *no-pair* terminology [111] references the fact that the projection operators effectively remove the pair creation and annihilation terms from the Dirac Hamiltonian.[9]

Intuitively one would expect that the various mathematical problems of the Dirac equation might disappear if the negative energy continuum states are removed. However, if the external field V_{ext} is the simple $1/r$ potential, corresponding to a point nucleus, $\hat{h}_+$ also has a critical charge at which the operator becomes w, just like the standard Dirac equation. The critical charge

$$Z_{cr} = \left(\frac{2}{\pi} + \frac{\pi}{2}\right)\alpha^{-1} \approx 124.16$$

is, however, lower in the projected case [114, 110].[10] This reinforces the idea that the issues related to the non-self-adjointness of the standard Dirac equation are related to the singular character of the $1/r$ potential, rather than the physical nature of the Dirac equation itself.

Unlike the Dirac equation, the no-pair operator has a lower bound in the sub-critical region, which is expected, as the negative continuum has been removed. That result was then further refined by Tix, showing that the operator's eigenvalues are in fact strictly positive [115, 116]. In general, however, the lowest eigenvalue is expected to be lower than the corresponding $1s_{1/2}$ bound state energy of the standard Dirac equation.

While there have been studies of the projected equation with a $1/r$ potential, the finite nucleus case seems to be unexplored so far. Like in the case of the standard Dirac equation, a singularity-free nuclear potential from a finite nucleus should remove the problem of non-self-adjointness and the critical charge. Futhermore, since the operator has no negative continuum, there should be no Z value where it becomes difficult to describe the bound states, as they never become degenerate with other states, and it should always be possible to obtain reasonable states and energies for any nuclear charge value, not unlike the non-relativistic Schrödinger equation.

Numerical implementation

When implementing the Hamiltonian (3.10) with FEM, it is not actually necessary to explicitly implement the projection operators. To explain this, let $|E(\boldsymbol{p}), \boldsymbol{p}, r\rangle$ and $|-E(\boldsymbol{p}), \boldsymbol{p}, r\rangle$ be the positive and negative energy continuum plane wave states of the free-particle Dirac Hamiltonian h_D, respectively. In other words, they obey

$$h_D \,|\pm E(\boldsymbol{p}), \boldsymbol{p}, r\rangle = \pm E(\boldsymbol{p}) \,|\pm E(\boldsymbol{p}), \boldsymbol{p}, r\rangle$$

Using those states, it is possible to formally construct a projection operator that restricts the Hilbert space to just the positive energy space with

$$\Lambda_+ = \sum_r \int \mathrm{d}\boldsymbol{p} \;\; |E(\boldsymbol{p}), \boldsymbol{p}, r\rangle \langle E(\boldsymbol{p}), \boldsymbol{p}, r|$$

[9] It should be pointed out that this Hamiltonian, while similar to, is not the same as the the no-pair Hamiltonian often used in relativistic quantum chemistry to avoid the continuum dissolution (also known as the Brown-Ravenhall disease) when electron-electron interaction operators (e.g. the Coulomb or Breit interaction operators) are introduced. In that case, the projection operator is usually constructed from the positive energy eigenstates of the full external field Dirac Hamiltonian, and so does not span quite the same space as the one based on the free-particle states. Furthermore, it is also dependent on the nuclear charge Z and the shape of the nuclear potential. [112, 113]

[10] It is also worth mentioning that the Dirac equation with a $1/r$ potential has another special nuclear charge value at $Z_* = (\sqrt{3}/2)\alpha^{-1} \approx 118.68$, where the operator loses self-adjointness if only normalizability is imposed as a boundary condition (due to a second solution to the equation becoming normalizable). The projected equation also exhibits the same special value, except at $Z_* = (3/4)\alpha^{-1} \approx 102.78$ [114, Eq. (2.9)].

Similarly, it would be possible to construct a projection operator for the negative energy subspace

$$\Lambda_- = \sum_r \int \mathrm{d}\boldsymbol{p} \ \ |-E(\boldsymbol{p}), \boldsymbol{p}, r\rangle \langle -E(\boldsymbol{p}), \boldsymbol{p}, r|$$

Either way, to restrict the Hamiltonian to one of the energy subspaces, it is only necessary construct the Hamiltonian (3.10) in the particular subspace as

$$h_\pm = \sum_{r,r'} \int\int \mathrm{d}\boldsymbol{p}\, \mathrm{d}\boldsymbol{p}' \, \langle \pm E(\boldsymbol{p}), \boldsymbol{p}, r| h_\mathrm{D} + V_\mathrm{ext} |\pm E(\boldsymbol{p}'), \boldsymbol{p}', r'\rangle \, |\pm E(\boldsymbol{p}), \boldsymbol{p}, r\rangle \langle \pm E(\boldsymbol{p}'), \boldsymbol{p}', r'|$$

While it would be possible to try to use the analytical free-particle solutions, in a basis-set approach, such as the finite element method, this would lead to linear-dependencies in the projected basis. Technically the projection operator should remove approximately half the basis functions, but when an analytical version is applied on a numerical basis, it does not actually remove any states.

FEM implementation

In the finite basis set approach, it is easy to approximate the projection operator by simply replacing all the states and operators with vectors and matrices. To obtain the (approximate) free-particle states numerically, the full Dirac equation can always be just diagonalized with $V_\mathrm{ext} = 0$, which yields a clear spectrum of positive and negative energy continuum states. The N_+ eigenstates that have positive eigenvalues ($E \geq mc^2$), could then be used to construct the approximate (N, N_+) matrix for Λ_+. This can then be used to construct the (N_+, N_+) matrix for the Hamiltonian (3.10).

The particular implementation in the code expands on the approach described in Section 3.4, and also solves the radial Dirac equation. As for the full equations, the projected Hamiltonian is also solved in each angular momentum subspace (labelled with the κ quantum number) separately, with the projection operator assumed to apply to each angular momentum block separately. In explaining the formalism, the first step is to split the full radial Dirac Hamiltonian matrix (3.8) in two, into first the free-particle part

$$T = \begin{pmatrix} m\alpha^{-2} S_\mathrm{LL} & -\alpha^{-1} D_\mathrm{SL} + \kappa\alpha^{-1} R^{(-1)}_\mathrm{SL} \\ \alpha^{-1} D_\mathrm{LS} + \kappa\alpha^{-1} R^{(-1)}_\mathrm{LS} & -m\alpha^{-2} S_\mathrm{SS} \end{pmatrix}$$

and then also the normalized nuclear potential part

$$\Phi = \begin{pmatrix} \Phi_\mathrm{LL} & 0 \\ 0 & \Phi_\mathrm{SS} \end{pmatrix}$$

The full matrix (3.8) can, of course, then be constructed by

$$H = T + Z\alpha\Phi$$

To determine the free particle states (in each angular momentum subspace), it is necessary to orthogonalize the T matrix via a Löwdin transformation $\tilde{T} = S^{1/2} T S^{-1/2}$, and then to diagonalize the orthogonalized matrix to obtain the states and energies

$$\tilde{T}\tilde{\phi}^0_i = E^0_i \tilde{\phi}^0_i \tag{3.11}$$

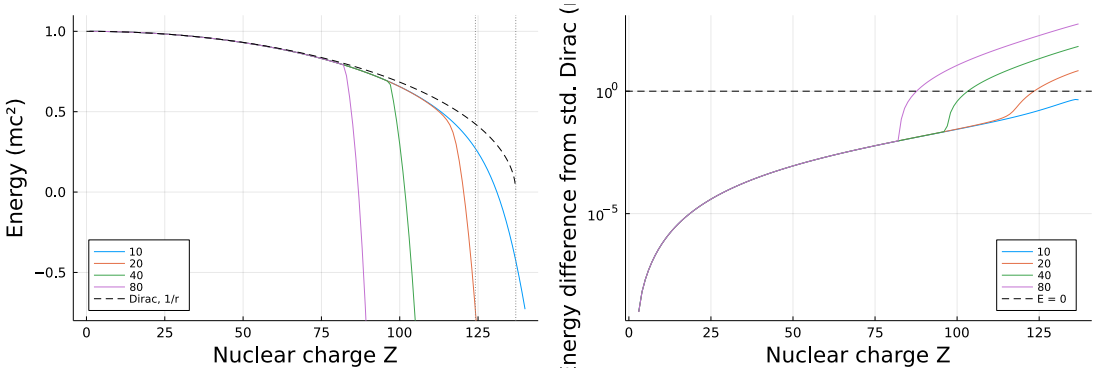

Figure 3.33: Numerical 1s ground state energy of the no-pair (projected) Dirac Hamiltonian for the point nucleus with a few different basis sets, differing by the number of elements. The figure on the left shows the absolute energy, whereas on the right we show the difference from the standard point-nucleus Dirac eigenenergy. The sudden turn in the curves is due to a non-physical numerical issue.

which approximately correspond to the free-particle plane wave states.[11] With those eigenstates, which clearly split into two distinct sets with $E_i^0 \geq mc^2$ or $E_i^0 \leq -mc^2$, is is possible to construct the projection operator corresponding to Λ_+

$$\tilde{P}^+_{:,i} = \tilde{\phi}^0_i \quad \text{for } i \text{ such that } E_i^0 \geq mc^2$$

where ":" indicates that the columns will be filled by the vectors corresponding to the eigenstates. It should be noted that this will not be a square matrix, and instead will usually have about half the columns, depending on the basis sets picked for the large and the small components. This matrix can then be applied on the Löwdin-orthogonalized Hamiltonian matrix $\tilde{H} = S^{-1/2} H S^{-1/2}$

$$H_+ = \tilde{P}^\dagger \tilde{H} \tilde{P}$$

to obtain the projected Hamiltonian matrix, corresponding to the projected operator h_+ of Equation (3.10). Diagonalizing H^+ yields the eigenstates and eigenenergies of the projected equation

$$H^+ \bar{\phi}^+_i = E^+_i \bar{\phi}^+_i$$

The matrix elements of H_+, and therefore the eigenstates $\bar{\phi}^+_i$, are not in the Löwdin-orthogonalized basis, but in the basis of the free-particle eigenstates of the T matrix. But the states can easily be transformed into the original FEM basis representation

$$\phi^+_i = S^{-\frac{1}{2}} \tilde{P} \bar{\phi}^+_i$$

which can then be used to plot the orbitals and so forth.

Point nucleus

While the point nucleus case, where the potential is simply

$$\hat{V}_{\text{ext}} = -\frac{Z}{r} \quad ,$$

[11] Technically, this can simply be done by setting $Z = 0$ in the solver of the Dirac equation of Section 3.4.

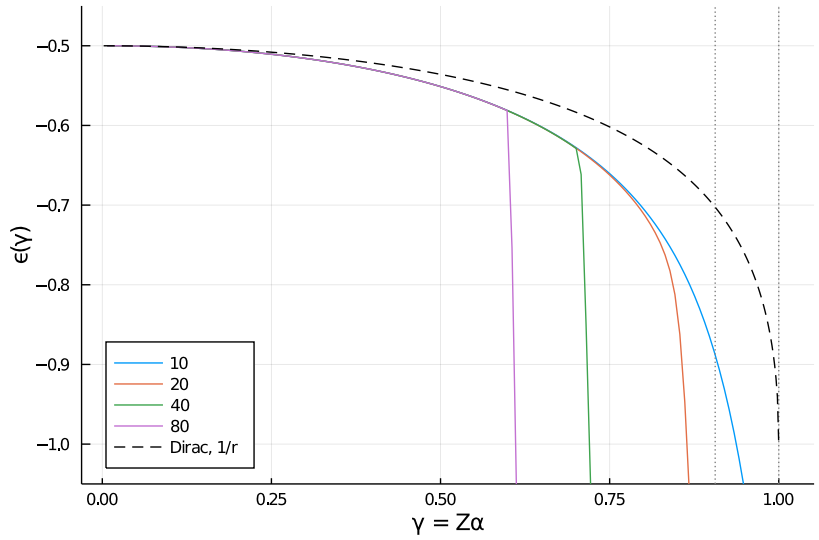

Figure 3.34: Same data as in Figure 3.33, but plotted on the same axes as FIG 1 of [114], showing qualitative agreement between the results, as long as the FEM approach does not run into numerical difficulties. The y- and x-axes are given by $\epsilon(\gamma) = (E - m)/\gamma^2 m$ and $\gamma = Z\alpha$, respectively.

is known to be problematic with a critical charge at $Z_{\text{cr}} \approx 124$ for the projected operator, it is still interesting to see how the method behaves. As can be seen in Figure 3.33, at smaller Z values, the method produces stable results. Generally, we see that the calculated energy is below the energy of the standard Dirac bound state energy, which is consistent with earlier results in the literature (e.g. when comparing Hardekopf's and Sucher's results [111] to Figure 3.34). However, at higher nuclear charge values, close to the critical charge, the numerical approach runs into issues. What is even more disturbing is that these numerical issues become worse as the basis is improved, which seems to imply that trying to converge by making the basis bigger and bigger seems to be problematic at any Z value.

Figure 3.35 shows that the issue is caused by a spurious state falling down through the spectrum, eventually becoming the lowest energy state. It should be emphasized that this appears to be a numerical error caused by the finite basis approach used, rather than a physical or a mathematical feature to the underlying equation. The plots on the left in Figures 3.36 and 3.37 illustrate the character of this spurious state. First, rather than having a clean, smooth bound state character, the state is highly oscillatory and, most importantly, has collapsed nearly to the origin. The only reason it has not collapsed completely into the origin is that the element closest to the origin has a finite size, and so there are limits to what states can be represented in that basis. This is clear from comparing the right plots between Figures 3.36 and 3.37 — the spurious state will always fill just the left-most element, and so is directly dependent on the choice of the element grid. Furthermore, visible on the bottom graphs of the same figures is that, in terms of the numerical free-particle basis, the spurious state is made up of the highest energy, most oscillatory free-particle solutions. This seems to imply that for a well-described bound state we need a balance, and the issue is that we run out of basis states to balance out the collapse into the centre.

To understand the balancing issue from a slightly different perspective, it is helpful to look at

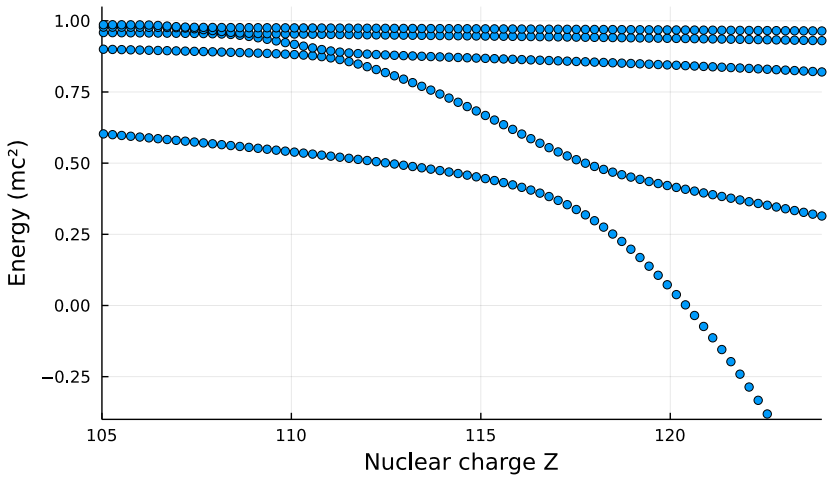

Figure 3.35: Low-lying states in the spectrum of the no-pair Hamiltonian for different Z values for a given basis. A spurious state falling through the spectrum is clearly visible.

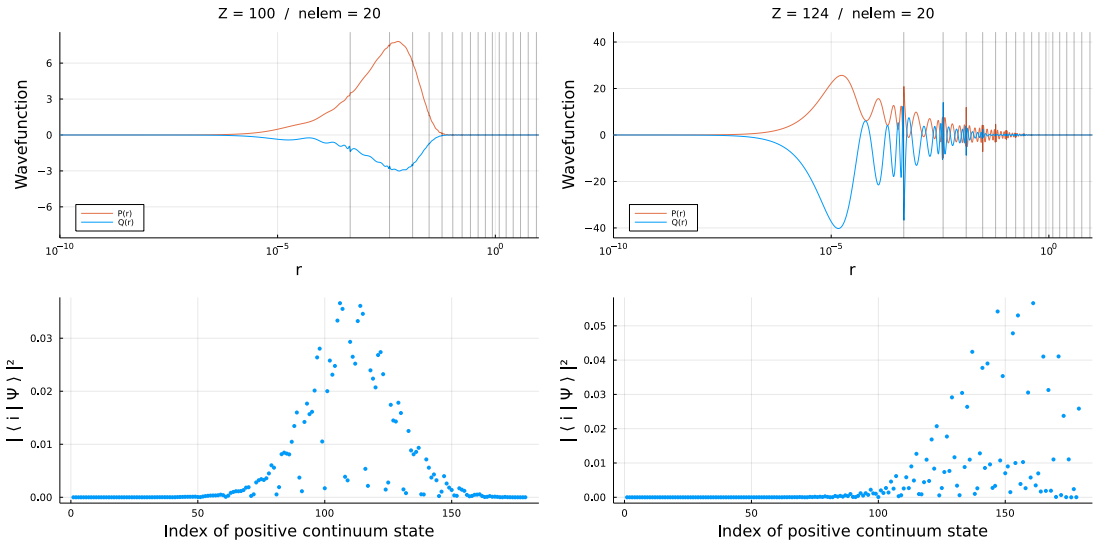

Figure 3.36: The FEM-based solutions for the 1s orbital of the no-pair Hamiltonian for a point nucleus at two different Z values (100 on the left, 124 on the right) for a relatively small basis (20 elements), illustrating the diving problem the numerical approach faces. The top graph shows the $P(r)$ and $Q(r)$ components of the orbitals. The bottom graphs show the magnitude squared of the coefficients in terms of the free-particle positive energy states used as a basis. Higher index implies a higher energy of the free-particle state.

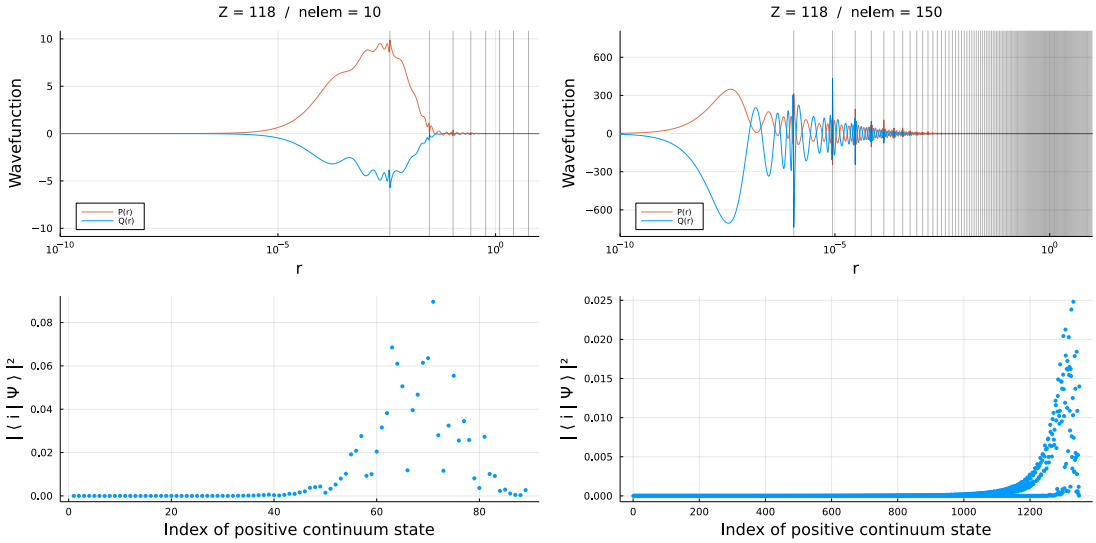

Figure 3.37: Same as Figure 3.36, but showing the orbitals and coefficients for different number of basis elements (10 on the left, 150 on the right) at $Z = 118$.

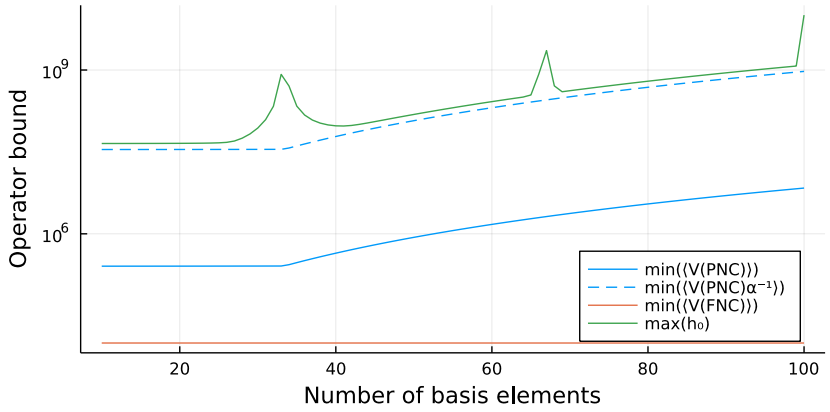

Figure 3.38: Bounds of the potential and free particle operators as a function of the number of elements in the FEM basis. The elements are distributed according to an exponential grid with an extra boundary at 10^{-4} a.u., corresponding to the radius of the *hollow nuclear model*, used to model the finite nucleus. The blue and orange lines are the absolute values of the lower bounds of the nuclear potentials for point (PNC) and finite (FNC) nuclei, respectively, with the dashed line scaled by α^{-1}, to represent the potential near the critical charge regions. The green line is the upper bound of the free-particle Dirac operator, with the kinks being caused by the grid boundaries from the exponential grid passing over the extra finite element boundary introduced for the hollow nuclear model.

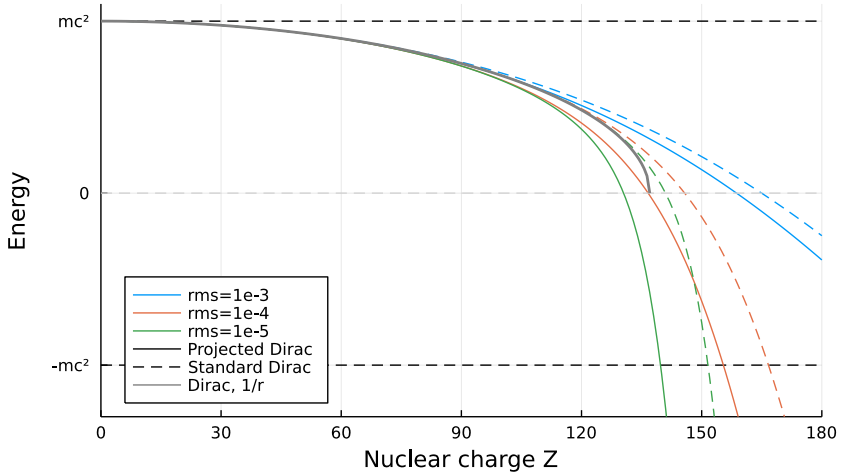

Figure 3.39: 1s energies for finite nuclear models in both the projected (solid) and unprojected (dashed) approach. The energies of the unprojected $1s_{1/2}$ states in the diving regime below $-mc^2$ are traced with the "Single-reference with energy filtering" strategy described in Section 3.5. The solid gray line represents the analytic Dirac $1s$ energy for a point nucleus.

how the bounds of the discretized operators change as the basis becomes bigger. This is shown in Figure 3.38, and it is very clear that the shape of the lower bound of the point nuclear potential and the upper bound of the free-particle operator are very similar, although the nuclear potential has a much smaller magnitude. However, the magnitudes become comparable near the critical charge, which is shown with a dashed line. This could be interpreted as a sort of a "tug-of-war" between the operators, with the free-particle Hamiltonian evidently losing out near the critical charge, but this needs further study, including in the non-relativistic limit.

Now, the full $1/r$ operator (i.e. when not represented in a finite basis) is, of course, unbounded. However, the mathematical analysis of the projected Hamiltonian shows that, before the critical charge, it is bounded from below but becomes unbounded after [110]. It looks as if the finite basis case is an approximate version of the issue that causes the full projected operator to break down. This does suggest a potential workaround though: it might be possible to bias our basis in a way to slightly prefer the free-particle operator so that the introduction of the spurious state could be avoided. However, in any case, in the case of the finite nucleus, the operator has a lower bound that is independent of the size of the basis, so that case should not run into these kinds of issues, and would hopefully converge well when the basis is enlargened.

Finite nucleus

As has already been mentioned, in the numerical FEM approach switching out the potential is trivial. The only thing that needs to be kept in mind is to make sure that any necessary extra element boundaries are in place at the discontinuities of the potential.

Figure 3.39 shows the calculated energies for the cutoff nucleus with various RMS values as a function of the nuclear charge Z. Like in the case of the point nucleus, the energy in the no-pair approach is always lower than the corresponding full Dirac energy. However, otherwise, the qualitative behaviour of the energy is very similar (e.g. larger nucleus leading to a higher energy).

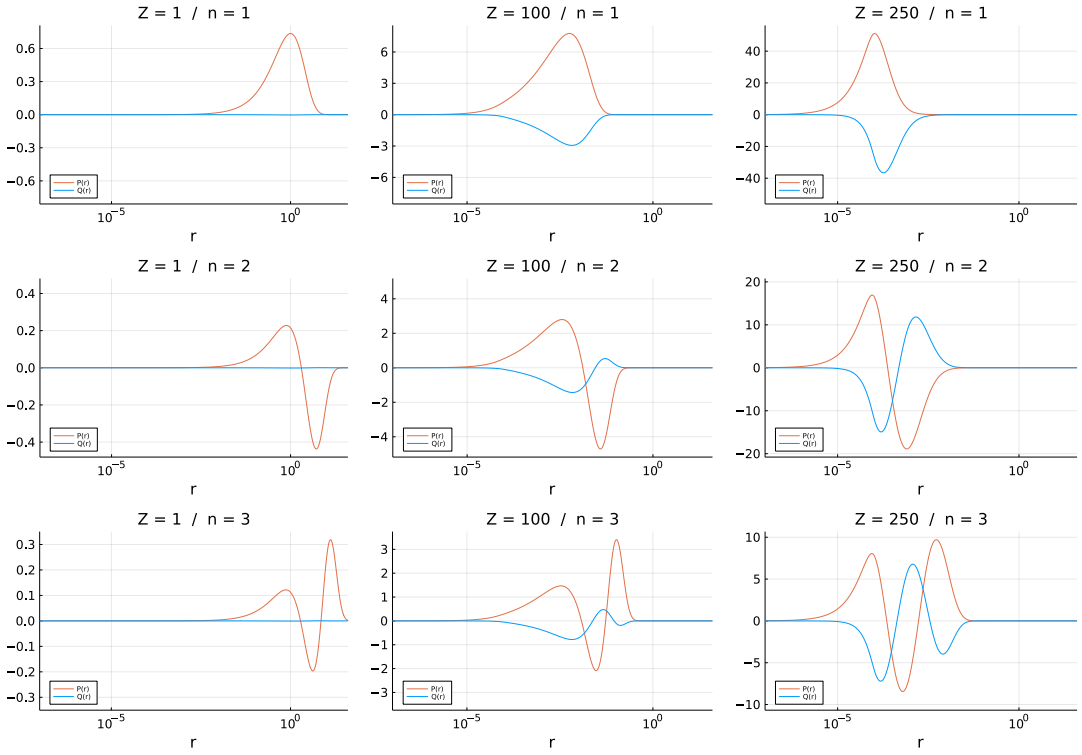

Figure 3.40: The s orbitals for a cutoff Coulomb potential ($r_{\text{rms}} = 10^{-4}$ a.u.) for a set of n (1–3, vertically) and Z (1, 100, 250; horizontally) values. The colors represent the $P(r)$ and $Q(r)$ functions, respectively.

One important thing to point out is that the method produces a smooth curve at energies lower than $-mc^2$. This does make sense since the negative energy states have been projected out and, therefore, there are no issues with degeneracies or resonances with the states of the negative energy continuum. Also, as was shown by Figure 3.38, the FNC potential operator is bounded from below and so does not introduce the numerical problems seen in the point nuclear case. Similarly, Figure 3.40 shows the $s_{1/2}$ orbitals for various n and Z values, and they are all smooth and behave as expected.

Another observation is that as Z gets larger, the small component $Q(r)$ becomes more and more significant, relative to the total density, visualized in Figure 3.41. It looks like that, roughly at the critical charge (i.e. when the energy becomes lower than $-mc^2$), there is kind of a transition from a predominantly large component regime to one where the small and large components have more or less the same density. This is to be expected if one considers the eigenstates of the $V(r)$ operator, which have symmetry between the P and Q component, so one would expect them to have exactly the same magnitude (this can be argued analytically from the form of the operator, but also has been verified numerically).

Finally, it is also important to show that the numerical results are actually reliable by checking the convergence of the energies. This is visualized in Figure 3.42, and it is clear that the method easily achieves systematic convergence to machine precision. Although, at higher Z values the required basis does get relatively large.

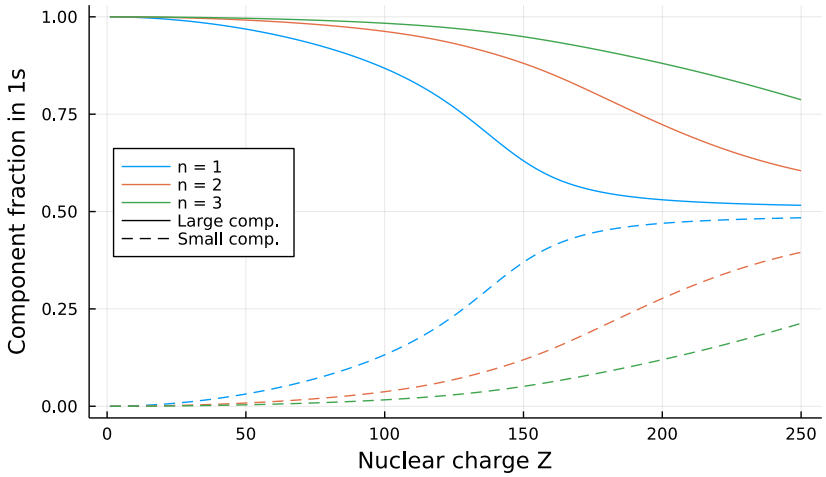

Figure 3.41: The fractions of large (solid) and small (dashed) components in the ns orbitals for $n = 1, 2, 3$, calculated as the density of the component (i.e. $\int |P(r)|^2 dr$ and $\int |Q(r)|^2 dr$).

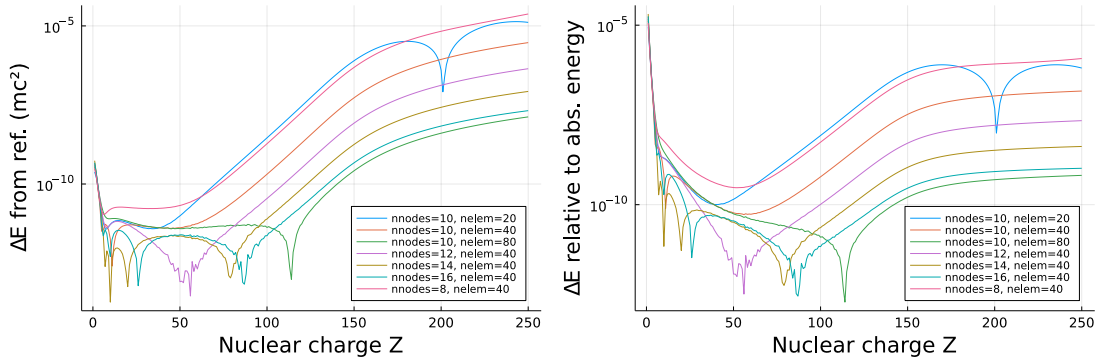

Figure 3.42: Convergence of the energy to the result of what we presume are the most accurate calculations with 160 elements. The graph on the left shows the absolute energy, which scales slightly with Z and the value on the right graph is relative to the energy $E - mc^2$.

Table 3.4: Comparison with experimental values (in eV) for the $1s$ ionization potential and a few transition energies of U^{91+}. The rows correspond to the standard hydrogenic Dirac equation with a point nucleus (PNC), finite nucleus with a *hollow nuclear model* (FNC), and the projected Dirac equation with the same finite nuclear model. The experimental values are taken from the 2019 review by Indelicato, by picking the values with the lowest uncertainty [2].

	Ionization potential		$1s_{1/2} \to 2p_{1/2}$		$1s_{1/2} \to 2p_{3/2}$	
	E	ΔE_{exp}	E	ΔE_{exp}	E	ΔE_{exp}
Dirac / PNC	132,279.93	454.83	98,064.45	458.84	102,630.10	451.98
Dirac / FNC	132,083.55	258.45	97,872.42	266.81	102,433.71	255.59
Projected / FNC	140,474.30	8,649.20	105,686.10	8,080.49	110,767.33	8,589.21
Experimental	$131,825.10 \pm 4.20$		$97,605.61 \pm 16.00$		$102,178.12 \pm 4.33$	

Experimental comparison: U^{91+}

What is clear so far is that the method seems to perform reliably for the projected calculations, at least for the finite nuclei. However, it would now be interesting to see how the theoretical results compare to the standard Dirac energies when pitted against experimental values. A good system for this is U^{91+}, which has been extensively studied both theoretically and experimentally [2] to test QED, and to calculate a few observables for it.

Table 3.4 compares the theoretical estimates (without any QED corrections) to experimental values. It is clear that using a finite nucleus noticeably improves the prediction from the Dirac equation. However, the predictions from the projected calculations are significantly worse than from the standard Dirac equation — with error increasing by more than an order of magnitude. This, unfortunately, does imply that, from a physical point of view, the projected Dirac Hamiltonian appears to be a worse starting point than the standard Dirac equation for further refinement of the states, such as with perturbation theory. That being said, it should nevertheless still be possible to use the projected states as the zeroth-order states in e.g. perturbative QED.

Interpretationally, however, these results lead to a bit of a conundrum. As discussed in Section 2.4 and further relying on the discussion in Sections 2.2 and 2.3, the eigenstates of the standard one-particle external field Dirac equation lead to superpositions of positive and negative energy free particle states. In the proper QED interpretation, which stems from the redefinition of the creation and annihilation operators of the negative energy states (to link them to positive-energy positrons) and therefore the vacuum state, the eigenstates of the Dirac equation are superpositions of one-electron and one-hole states. The projected Hamiltonian (3.10) disallows such superpositions by construction, and therefore the resulting states are the seemingly more physical, pure-electron states. Naively, one could expect that the resulting energies would also correspond more closely, or at least as closely, to the experimental energies. As evident from the numerical results though, that is not the case.

When thinking about the Dirac Hamiltonian with an interaction term (2.15) in the full Fock space QED picture, the projection operators have the effect of removing the pair creation and annihilation terms. Reintroducing them should therefore recover the predictions of the Dirac equation. However, what needs to be remembered is that when linking the single-particle picture to the Fock space picture, there are two choices for the reference vacuum states: (1) the non-empty Dirac sea vacuum $|\tilde{0}\rangle$ filled with positronic states, and (2) and the true vacuum $|0\rangle$ void of all particles. For the projected Hamiltonian, which only acts on the one-electron subspace, there is no difference, but the choice matters when the pair creation and annihilation operators are reintroduced. If c_i and d_i are the annihilation operators corresponding to the discretized free particle states from Equation (3.11), then the solutions of the standard Dirac equation are given

by the superposition

$$|\Psi\rangle = \sum_i \Psi_i^{(1,\infty)} c_i^\dagger \,|\tilde{0}\rangle + \sum_i \Psi_i^{(0,\infty-1)} d_i \,|\tilde{0}\rangle \tag{3.12}$$

coefficients $\Psi^{(n,m)}$ represent the expansion coefficients in the (n,m)-particle subspace, which in the case of the $|\tilde{0}\rangle$ vacuum state has an infinite number of positrons. However, in the proper QED reinterpretation of the operators, the solution of the full Hamiltonian with the creation and annihilation operators would become an infinite sum

$$|\Psi\rangle = \sum_i \Psi_i^{(1,0)} c_i^\dagger \,|0\rangle + \sum_{ijk} \Psi_{ijk}^{(2,1)} c_i^\dagger c_j^\dagger d_k^\dagger \,|0\rangle + \dots \tag{3.13}$$

mixing in all the (n,m)-particle subspaces that have the same total electric charge.[12] It should be noted that the higher particle number states come with a severe energy penalty, as each additional electron-positron pair brings $+2mc^2$ of mass energy.

The cases (3.12) and (3.13) are unlikely to be equivalent, and hence can not yield the identical results. But it is expected that including the extra determinants in (3.13) should reduce the discrepancies in the observables shown in Table 3.4. One thing that should be noted is that, in the latter case, the Hamiltonian is bound from below and, therefore, the variational principle would apply: increasing the size of the vector space would lower absolute energies. While this would mean that the absolute energies, which are already lower for the projected case, would move even farther from the standard Dirac equation solutions, this does not mean that the differences in energies, which are the observable quantities, could not improve. Relatedly, the zero-charge ground state, which in the free case is just the vacuum $|0\rangle$, would also be a superposition of states with various particle numbers

$$|\Psi\rangle = \Psi_i^{(0,0)} \,|0\rangle + \sum_{ik} \Psi_{ik}^{(1,1)} c_i^\dagger d_k^\dagger \,|0\rangle + \sum_{ijkl} \Psi_{ijkl}^{(2,2)} c_i^\dagger c_j^\dagger d_k^\dagger d_l^\dagger \,|0\rangle + \dots \tag{3.14}$$

This also has the implication that its energy would be lower than zero, which in turn would have implications for the calculation of e.g. ionizations potentials. In the FEM approach, it would be relatively straightforward to implement a numerical solver to determine the wavefunction (3.13) up to a certain order, and it would be interesting to see the resulting energies and how they compare to the experiment.

3.7 Implementation & code

The equations described in this chapter were all solved using a finite element basis set for the radial components of the single-particle functions [88]. Fortunately, it was not necessary to reimplement the basis set, as the code underlying the referred paper was well-organized and published on GitHub under an open-source license. However, in order to make the development easier and more efficient, the goal was to be able to implement the physical equations in the Julia programming language [100, 101]. This was desirable because it was clear that the work will be very exploratory in nature, meaning that the implementation would continuously have to be modified. As such, using a dynamic language with a REPL-type workflow helps considerably, while Julia allows the code to still to compile to machine code and run as fast as possible. This section briefly describes some of the technical details underlying the programming part of this work.

[12] For completeness, it should be noted that not all the index combinations for the coefficients $\Psi^{(n,m)}$ are allowed, as one has to take into account the Pauli principle.

Reorganizing the HelFEM C++ code

In order to utilize the existing HelFEM code in Julia, it was necessary to partially reorganize the code. While the code was already organized in a logical, modular way, an issue with the original implementation was that it was oriented towards compiling executables from the same source code. The code also had external dependencies, some relatively heavy and not necessary for the finite element method part.

As part of this work, many of the C++ classes and functions were moved into a separate `libhelfem` library. This included the code necessary to represent the polynomial bases, the radial bases and to evaluate the matrix elements of operators, including of the various nuclear models. This part of the code now compiles separately into a static library and only depends on the Armadillo and GSL libraries. All the modifications have been upstreamed into the official HelFEM Git repository[13].

In the future, it would be beneficial to continue refactoring the code in this vein. It would be ideal to have a high-quality C or C++ library available that can evaluate the finite element basis functions and matrix elements, with bindings to other languages such as Fortran, Python, and, of course, Julia.

Julia interface for HelFEM

The Julia interface to the HelFEM C++ code is via a Julia package, `HelFEM.jl`.[14] The package exposes the various features of the HelFEM code as Julia types

PolynomialBasis Wraps the corresponding `PolynomialBasis` class from the C++ code, and represents a set of polynomials on an interval $[-1, 1]$.

RadialBasis Takes a `PolynomialBasis` as an argument and constructs the corresponding `RadialBasis` object, which represents a full finite element on some set of elements.

FEMBasis Similar to `RadialBasis`, except this offers a pure-Julia implementation of the part that generalized a `PolynomialBasis` to a finite element basis across many elements. The reimplementation was done to improve on a convention chosen in `RadialBasis`. Unlike `RadialBasis`, this object can actually be used to represent basis sets on any interval $[a, b]$, whereas `RadialBasis` only supports $[0, r_{\max}]$.

Each of these types can be interacted with the help of various functions and methods defined in the package. For full documentation, see the developer documentation of `HelFEM.jl`.[15]

In order to expose the existing HelFEM C++ classes and methods to Julia, it was necessary to wrap it using the `CxxWrap.jl` package[16]. This exposes the necessary Julia bindings as a shared library, which is distributed as a binary artefact available via the generated `HelFEM_jll` wrapper package.

Implementation of physical equations

The implementation of the physical Schrödinger and Dirac equations are not currently organized into a package. However, the implementation on top of a `FEMBasis` object follows almost directly

[13] Available on GitHub: `https://github.com/susilehtola/HelFEM`

[14] Currently unregistered and only available via the GitHub repository: `https://github.com/mortenpi/HelFEM.jl`

[15] Available at `http://mortenpi.eu/HelFEM.jl/dev/`

[16] Available via the Julia package manager and the source code at `https://github.com/JuliaInterop/CxxWrap.jl`.

the formalism described in each of the sections of this chapter. The code provides the necessary methods to evaluate the matrix elements of all component operators, which are then simply added together to form the Hamiltonian. The Hamiltonian is then diagonalized with Julia's built-in eigenvalue solver.

A future plan of development is to generalize the implementation of physical operators to use basis sets implemented using the `ContinuumArrays` interface. It provides a general basis-agnostic way of describing physical or mathematical problems and would allow other basis sets to be used, such as the ones implemented in `CompactBases`.[17] On the other hand, implementing the `ContinuumArrays` interface to `HelFEM.jl` would also allow HelFEM to be easily used in other existing projects that are built upon the general framework.

[17]Both packages are available via the Julia package manager, with the source code hosted on GitHub at `https://github.com/JuliaApproximation/ContinuumArrays.jl` and `https://github.com/JuliaApproximation/CompactBases.jl`.

Chapter 4

Relativistic many-particle theory and QED

While Chapter 3 dealt with the single-particle Dirac equation, the focus in this chapter is on the many-electron case. As discussed in Section 2.5, accounting for the QED effects in many-particle systems is a non-trivial matter, especially for large systems (i.e. the tail end of the periodic table). This chapter describes the work that went into expanding the treatment of QED effects in the 2018 version of the GRASP software [22].

GRASP — which stands for General Relativistic Atomic Structure Package — has a long history in atomic physics, going all the way back to 1970s [117, 98, 118, 119, 120, 121, 22]. It is a software package written in the Fortran programming language for relativistic many-electron atomic calculations. At its core, it uses the multiconfigurational Dirac-Hartree-Fock (MCDHF) approach to determine the energies and many-electron wavefunctions of the electrons in atoms. Those results can then be further improved and elaborated upon via additional configuration interaction (CI) calculations and by running additional programs that are able to calculate various other properties, such as multipole transition properties, spectral line widths, isotope shifts [122] or the hyperfine energy splittings due to the Zeeman effect [123, 124].

GRASP already contains an implementation of the QED self-energy and vacuum polarization corrections in the CI portion of the software. However, the self-energy contribution specifically is implemented in a very simplistic manner, as a perturbative corrective energy shift derived from tabulated values computed for hydrogen-like orbitals. Since the original implementation in GRASP, various effective operators have been proposed as alternatives for including self-energy in a many-electron calculation [18, 19, 20, 73], some of which have already been implemented and employed in many-electron calculations with other versions of GRASP [125]. The core goal of this work, therefore, is to expand and generalize the QED self-energy implementation with the new effective operators, so that users of GRASP could estimate QED effects more reliably at the correlated level.

Finally, a note on the literature. While the following sections do give a concise overview of the necessary theory to understand the implementation of the QED operators, there are other resources out there that offer a more comprehensive exposition.

- The 2007 book by Ian Grant could in many ways be considered to be the theory manual for the GRASP software, focussing specifically on the relativistic theory and covering more or less every aspect of the theory [62].
- The non-relativistic multiconfigurational Hartree-Fock (MCHF) method is presented in the 1997 book by Fischer, Brage, and Jönsson [126]. A 2016 review by the same authors also

concisely covers the background, including relativistic methods [127].

- Also, while not directly relevant for this work, the 2007 book by Dyall and Faegri is an excellent reference on relativistic quantum chemistry, including the various transformations of the Dirac Hamiltonian to two-component forms [113].

4.1 Many-particle atomic structure theory

When formulating wavefunction-centric approaches[1], such as the MCDHF and CI methods central to GRASP, it is necessary to have a precise definition of how such wavefunctions are represented. While it definitely is possible to represent N-particle wavefunctions with, for example, a $3N$-variable function $\Psi(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N)$, the most general and convenient approach is still to start from a Fock-space formalism[2], briefly summarized in Section 2.2. This is even more true when it becomes necessary to take into account the symmetries of the wavefunction due to the inherent indistinguishability property of elementary particles of the same species.[3] Furthermore, as it is not really possible to numerically represent full QED wavefunctions (neither the photon component nor pair creation; see Section 2.5), such wavefunctions only contain the (multi-)electron component, and interactions beyond the electronic component are taken into account via effective operators.

Many-particle basis and configuration state functions

In the Fock-space formalism, defining many-particle states depends on having single-particle basis states that are then used to construct the many-particle states via tensor products. As an atomic Hamiltonian usually has rotational symmetry and therefore commutes with the total angular momentum operators $\boldsymbol{J} = \boldsymbol{L} + \boldsymbol{S}$, it makes sense to take advantage of that. The single-particle states, usually referred to as *orbitals*, are therefore generally assumed to be eigenstates of total angular momentum (see also Section 2.6).

In GRASP, the orbitals $|n\kappa m\rangle$ are defined in the same way as they were defined in Eq. (2.28) when analysing the radial Dirac equation, i.e.

$$|n\kappa m\rangle = \psi(r, \theta, \varphi) = \frac{1}{r} \begin{pmatrix} P_{n\kappa}(r)\chi_{\kappa m}(\theta, \varphi) \\ iQ_{n\kappa}(r)\chi_{-\kappa m}(\theta, \varphi) \end{pmatrix} \tag{4.1}$$

where $P_{n\kappa}(r)$ and $Q_{n\kappa}(r)$ are some radial functions, and $\chi_{\kappa m}(\Omega)$ are the spin-spherical harmonics. Each of these orbitals is then labelled first by their angular momentum quantum numbers κ (or,

[1] As opposed to methods that focus on calculating observables or eigenvalues directly, sidestepping the wavefunctions themselves, such as perturbation theory. That said, in such methods, the underlying theory is still based on a Fock-space formalism, and it is just not as important to explicitly think about that. Rather, one often operates on a higher level of abstraction, such as considering Feynman diagrams and writing down the necessary integrals directly from those.

[2] Interestingly, in a formalism where wavefunctions are just multivariable functions, there would technically be a separate, distinct mathematical model and Hamiltonian for each particle number, as the function space and the number of terms in the Hamiltonian are different in each case. Furthermore, it is not clear how one would create superpositions of states with differing particle numbers, in the circumstances where that might be necessary. This is in strong contrast to the Fock-space approach, where states with any number of particles can be described within the same mathematical model.

[3] According to the spin-statistics theorem, many-fermion wavefunctions are anti-symmetric under the interchange of particle coordinates, whereas many-boson wavefunction are symmetric. As a more concrete example, for a two-particle wavefunction, the possible wavefunctions are constrained by $\Psi(\boldsymbol{x}_1, \boldsymbol{x}_2) = \pm\Psi(\boldsymbol{x}_2, \boldsymbol{x}_1)$, with $+$ or $-$ for bosons or fermions, respectively.

equivalently, ℓ and j), and then by their principal quantum number n within each angular momentum block.

As the orbitals simply form as a basis, then, in principle, there is no one specific prescription for choosing the radial function $P_{n\kappa}(r)$ and $Q_{n\kappa}(r)$. They might be something as simple as the radial solutions of the radial Dirac equation (2.30) with a given potential (e.g. $1/r$). Or they could be defined via another basis set for the radial functions themselves as was done in Chapter 3. In GRASP, the orbitals are recalculated over and over again in a multiconfiguration Dirac-Hartree-Fock procedure [98, 127, 62, 126]. In a nutshell, it solves the radial Dirac equation, but the equation also includes an effective average Coulomb potential from all the electrons of the system, providing a mean-field approximation to the inter-electron interaction.

Something that should be mentioned is that such orbitals would, in general, mix electron and hole states, in the sense discussed in Section 2.4. While this may create interpretational questions regarding the many-electron wavefunction, this does not create any inherent mathematical issues, and the methods do produce good, reliable results when compared to experimental data.

Many-electron basis

Each of the aforementioned orbitals is associated with a creation operator $c_i^\dagger$ and annihilation operator c_i. These can be used to write down the many-particle product basis states

$$|i_1 i_2 \dots i_N\rangle = c^\dagger(i_1)c^\dagger(i_2)\cdots c^\dagger(i_N)\,|0\rangle$$

In principle, by iterating over all the combinations of the single-particle states, all the while taking into account the Pauli principle, it is possible to generate a complete N-particle basis corresponding to this particular single-particle basis. In case it is necessary to convert to a multi-variable function representation of the fermionic N-particle states, each of basis states $|\Psi\rangle = c_1^\dagger c_2^\dagger \cdots c_N^\dagger\,|0\rangle$ can be written as a determinant

$$\Psi(\boldsymbol{x}_1,\dots,\boldsymbol{x}_N) = \frac{1}{\sqrt{N!}}\begin{vmatrix} \psi_1(\boldsymbol{x}_1) & \psi_2(\boldsymbol{x}_1) & \cdots & \psi_N(\boldsymbol{x}_1) \\ \psi_1(\boldsymbol{x}_2) & \psi_2(\boldsymbol{x}_2) & \cdots & \psi_N(\boldsymbol{x}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \psi_1(\boldsymbol{x}_N) & \psi_2(\boldsymbol{x}_N) & \cdots & \psi_N(\boldsymbol{x}_N) \end{vmatrix}$$

and it is the matrix determinant that ensures the anti-symmetry of the resulting multivariable function.[4] These functions are known as *Slater determinants* [128], and it is common to use that term interchangeably when discussing fermionic many-particle product states.

It should be emphasised here that the coordinate vector $\boldsymbol{x}$ can formally consist of more than just the standard 3-vector coordinates of space. It would also include any additional coordinates and labels of the single-particle wavefunction would have, such as the label of the component of a 4-component Dirac wavefunction.

Configuration state functions

It would be possible to represent the many-electron wavefunctions as superpositions of Slater determinants, and some many-electron programs do opt for that [129, 103, 93]. However, this has two major disadvantages: (1) the combinatorial explosion means that the basis becomes very large very quickly, and (2) the Slater determinants, even though products of orbitals that

[4] There is a bosonic equivalent for the Slater determinant as well, which can be written as a *matrix permanent* of a similarly organised matrix of single-particle functions.

are angular momentum eigenstates themselves (i.e. eigenstates of $\boldsymbol{J}^2$ and J_z), are not usually eigenstates of total angular momentum $\boldsymbol{J}^2$. Using configuration state functions (CSFs) that have been symmetrised for angular momentum can be used to work around both of these issues. Conceptually, this is very similar to constructing total angular momentum eigenstates of products of two angular momentum eigenstates via Clebsch-Gordan coefficients (see Section 2.6), essentially finding a unitary transformation of the full set of Slater determinants into a set where all the states are angular momentum eigenstates. As this is a unitary transformation, the set still spans the exact same space as the original set of Slater determinants and forms an orthonormal basis, so diagonalising the Hamiltonian leads to the exact same answer, simply in a different representation.

The benefit of CSFs is that it is only necessary to consider a subset of them that have the desired total angular momentum and parity, as a rotationally invariant Hamiltonian is block diagonal in terms of total angular momentum J, the M quantum number, and parity. As the angular momentum of the target state is generally known, most of the CSFs can be ignored. This is in contrast to Slater determinants, where it is necessary to consider a much larger set of objects, as they can also be coupled into states with other total angular momenta. Overall, this leads to a large reduction in the number of many-particle basis elements. In addition, it also has the effect of forcing the resulting state to have the correct angular momentum in the numerical procedure, which can be useful in variational procedures like Hartree-Fock.

For a more precise definition of a CSF, it is first necessary to define two other terms:

Subshell A set of orbitals with fixed principal quantum number n and total angular momentum quantum number κ. It is a set of $2j + 1$ states, with all the different possible values of the m quantum number, e.g. a $2s_{1/2}$ or $5f_{5/2}$ subshell.

Configuration A collection of orbitals is described by a list of subshells and their *occupation numbers* — the latter meaning how many times each subshell is represented. The m quantum numbers are not specified, and rather it represents a set of all such Slater determinants that correspond to the given subshells and occupation numbers, e.g. a $1s^2_{1/2}2s^2_{1/2}2p^1_{1/2}$ configuration.

A CSF is the superposition

$$|C; J, M, \gamma\rangle = \sum_{\boldsymbol{m}} C(C, J, M, \gamma; \boldsymbol{m}) \, |C; \boldsymbol{m}\rangle$$

of the Slater determinants $|C; \boldsymbol{m}\rangle$ of a configuration C, where $\boldsymbol{m}$ represents the m quantum numbers of all the orbitals. The coefficients $C(C, J, M, \gamma; \boldsymbol{m})$ are analogous to the Clebsch-Gordan coefficients and are fully determined by the angular momentum algebra, and hence completely independent of the radial parts of the orbitals.

The additional label γ is necessary to distinguish between states that have the same total angular momentum quantum numbers but are actually different states. This is something that becomes an issue when coupling more than two states together. For the most part, in practice, it means that one has to specify the order in which subshells are coupled together, leading to a coupling tree, but it is also an issue within subshells with more particles. In the latter case, states are systematically often distinguished using Racah's seniority quantum number, although this scheme also breaks at higher j-values and numbers of particles [130, 131, 132, 133, 134].

Determining the angular coefficients $C(C, J, M, \gamma; \boldsymbol{m})$, however, is a non-trivial problem. A further complication, due to the anti-symmetry requirements of the product states, is that certain product states are not allowed when coupling states with the same n and κ values (i.e. when putting multiple electrons on the same orbital). It has the practical effect that not all total

angular momentum values J satisfying the *triangle relation*, familiar from the Clebsch-Gordan coefficients, are allowed. In principle, it is possible to numerically diagonalize the $\boldsymbol{J}^2$ operator [135]. However, more sophisticated approaches exist to do this more efficiently, such as the methods of fractional parentage and grandparentage [76, 136]. In GRASP, a substantial amount of work has gone into making this procedure as efficient as possible [137, 138, 139].

Operators

In the Fock-space formalism, operators such as the Hamiltonian of the physical system can be represented in terms of sums of products of creation an annihilation operators. The coefficients of these expansions are the matrix elements of these operators.[5] When working in the effective Hamiltonian approach, there are two general classes of operators of interest.

Single-particle operators. These operators act on the different particles individually, such as the free Dirac operators or a nuclear potential, and can be written as

$$T = \sum_{ij} T_{ij}\ c_i^\dagger c_j$$

In an orthonormal basis, the matrix elements T_{ij} can be calculated via

$$T_{ij} = \langle i|T|j\rangle = \langle 0|c_i T c_j^\dagger|0\rangle = \int \mathrm{d}\boldsymbol{x}\ \psi_i^\dagger(\boldsymbol{x})\hat{T}\psi_j(\boldsymbol{x})$$

where $\hat{T}$ is the equivalent operator acting in the function representation. In the functional representation, such operators act on just one set of coordinates.

Two-particle operators. These operators represent interactions between particles and can be represented as

$$V = \sum_{i<j,k<l} V_{ijkl}\ c_i^\dagger c_j^\dagger c_k c_l$$

Examples such operators are the effective electron-electron interaction operators, e.g. the Coulomb or Breit interaction. Due to the fermionic symmetry, the matrix elements have restrictions, and one way to take that into account is simply to disallow certain indices. Alternatively, one can also simply impose

$$V_{ijkl} = -V_{jikl} = -V_{ijlk} = V_{jilk}$$

which also leads to any matrix elements having repeating ij or kl pairs to be zero. In that case, the expansion can be summed over all possible index combinations

$$V = \frac{1}{4}\sum_{ijkl} V_{ijkl}\ c_i^\dagger c_j^\dagger c_k c_l$$

The matrix elements can be calculated via double integrals

$$V_{ijkl} = \langle ij|B|kl\rangle = \langle 0|c_i c_j V c_k^\dagger c_l^\dagger|0\rangle = \int\int \mathrm{d}\boldsymbol{x}_1\,\mathrm{d}\boldsymbol{x}_2\ \psi_i^\dagger(\boldsymbol{x}_1)\psi_j^\dagger(\boldsymbol{x}_2)\hat{V}\psi_k(\boldsymbol{x}_1)\psi_l(\boldsymbol{x}_2)$$

as long as the single-particle states form an orthonormal basis.

[5] It should also be noted that, arguably in a more physical way, the operators can also be written directly in terms of field operators.

It is interesting to note here that the total angular momentum operator $\boldsymbol{J}^2$, when thinking in this many-particle Fock-space representation, is not a pure one-particle operator but also contains a two-particle component. This insight is generally lost when angular momentum operators are discussed solely in the context of single-particle states.

Determining stationary states

The central task of a software like GRASP is to determine the multi-electron wavefunctions of the stationary states of an atom. Or, in other words, the eigenstates and their corresponding eigenvalues of the Hamiltonian H. The eigenstates are assumed to be superpositions of the previously defined CSFs

$$|\Psi; J, M\rangle = \sum_{C,\gamma} \Psi_n \, |C; J, M, \gamma\rangle \tag{4.2}$$

where the sum runs over all possible configurations C and couplings γ within those configurations. In principle, once the eigenstates have been determined, it is relatively straightforward to calculate any additional observables or other quantities that might be of interest.

There are a few well-known options for solving many-electron systems in this framework, and each come with their own pros and cons:

Exact diagonalisation Exact diagonalisation would be the most naive way of determining the eigenstates in this approach. Given a large enough set of single-particle states, one can generate the corresponding set of multi-electron states (Slater determinants or CSFs), calculate all the matrix elements of the Hamiltonian, and then simply diagonalise the Hamiltonian matrix in this basis. This would yield a good answer, but the computational cost here is far too large, once the system has more than a handful of electrons. As such, it is necessary to employ some approximate scheme.[6]

Another issue with this approach is that there is no *a priori* obvious way to pick a good set of orbitals. In general, it is necessary to span the single-particle space, which usually means a large set of orbitals, in turn exacerbating the computational cost problem.

Mean-field or Hartree-Fock method In this case, a single configuration is picked and is used to construct the corresponding CSF or Slater determinant $|\Psi\rangle$. Then, rather than working with the expansion coefficients of the wave function expansion, the goal is to optimise the radial parts of the orbitals such that the energy expectation value of the Hamiltonian

$$\langle \Psi | H | \Psi \rangle$$

is minimised. While it is possible to use more sophisticated optimisation algorithms, often the orbitals are updated in a simple iterative manner, by starting from a guess for the orbitals, diagonalising the Hamiltonian and repeating, with each iteration then stepping closer and closer to a converged state.

One key drawback of the Hartree-Fock approach is that the resulting wavefunction will, by definition, always be described by a single determinant or CSF. This omits the *static and dynamic correlation* contribution, a form of quantum entanglement of the electrons, arising from the extra degrees of freedom that are introduced when including additional configurations.

[6] In principle, there are methods for performing exact diagonalisation of large systems via stochastic methods, such as the Full CI Quantum Monte Carlo (FCIQMC) approach [140].

Configuration interaction (CI) A way to improve on the mean-field approach is to, essentially, step towards exact diagonalisation in a step-by-step fashion, systematically including more and more additional CSFs in your calculation.[7] In CI, this is done as a separate calculation after Hartree-Fock, where the Hartree-Fock optimised orbitals, used as the basis for the single-particle space, are kept fixed.

This approach works in many cases because many of the expansion coefficients of the CSFs in expansion (4.2) will have very small or even negligible magnitudes. This is helped along by having Hartree-Fock optimised orbitals as the single-particle space, which helps to bias the magnitudes of the coefficients towards certain CSFs. As such, many of the additional CSFs can be omitted without introducing large errors, and there are schemes to help choose the most approriate CSFs. However, like the exact diagonalisation approach, the many-particle space still becomes very large very quickly, and so it is often not feasible to achieve numerical convergence.

Multiconfigurational Hartree-Fock (MCHF) and multireference CI (MRCI) It is possible to extend both Hartee-Fock and CI by using multiple reference determinants or CSFs. The multiconfiguration Dirac-Hartree-Fock (MCDHF; i.e. relativistic MCHF) method implemented in GRASP does just that, where it is possible to optimise a superposition of multiple CSFs by both varying the orbitals and also the coefficients of the superposition. This allows a large part of the static correlation to be taken into account in the variational step, which helps a lot in open-shell systems, for example.

Optimising the orbitals for superpositions of large sets of determinants or CSFs is computationally expensive, however. Therefore it is still necessary to employ a CI method to further refine the wavefunction, but in this case the CI space will be generated from multiple reference states.

The Hamiltonians in GRASP

The core part of the GRASP implements the relativistic version of the MCHF method based on the Dirac equation (multiconfigurational Dirac-Hartree-Fock, or MCDHF). In the MCDHF step, GRASP variationally optimises the orbitals under the *Dirac-Coulomb Hamiltonian*

$$H_{\mathrm{MCDHF}} = h_{\mathrm{Dirac}} + V_{\mathrm{Coulomb}}$$

where h_{Dirac} is the singe-particle Dirac Hamiltonian with a nuclear potential and V_{Coulomb} refers to the $1/r$ Coulomb interaction between the electrons. For the nuclear potential, GRASP supports the standard $1/r$ potential from a point nucleus, but more commonly the two-parameter Fermi model is used with parameters that reproduce the physical RMS of the isotope under study.

Once the orbitals have been optimised in the MCDHF step, the many-electron wavefunctions can be further refined in a relativistic multireference CI procedure. A major benefit of the follow-up CI step is that it allows for the many-particle space (i.e. the list of CSFs included in the calculations) to be expanded considerably. This is because optimizing the orbitals in the MCDHF step is a much more expensive operation, compared to just the relatively simple construction and diagonalization of a many-particle Hamiltonian matrix performed in CI. That

[7] This process is generally referred to as *substitution*. The resulting CSFs are usually classified by the number of electrons that have been moved relative to the reference CSFs. Single (S) and double (D) substitutions are generally sufficient to describe most of the correlation in a system. This can also be justified through Z-dependent perturbation theory [126].

said, generally, this means that the many-particle space gets expanded to be as large as possible, so it is still an expensive step, comparable to the MCDHF calculation itself.

In GRASP, the radial components $P_{n\kappa}(r)$ and $Q_{n\kappa}(r)$ of the orbitals are represented on a numerical grid, as opposed to using an analytical basis set. Hence the Dirac-Hartree-Fock equations are solved for each orbital separately using numerical integration, with what is essentially the *shooting method* [98]. However, unlike in a diagonalisation with a basis set, this does not leave "leftover" orbitals that can be used as *virtual orbitals*, that are necessary for generating the many-particle configurations via substitution. Instead, the procedure for generating virtual orbitals involves adding in new configurations that occupy the desired orbitals and then re-run the MCDHF calculation. In practice, this is done in a so-called *layer-by-layer* fashion, where only the orbitals with the next highest principal quantum number n are optimised at once, with the already optimised orbitals kept fixed [127]. This is mainly due to the difficulties GRASP has when converging the orbitals when too many are allowed to be varied simultaneously.

The CI program in GRASP implements additional physical corrections which can be included as extra terms in the Hamiltonian:

Breit interaction. Breit is the most significant of the corrections in terms of its effect on, for example, the energy eigenvalues. GRASP, in fact, implements the frequency-dependent version of Breit, of the form (2.20) [117]. However, there are interpretational difficulties with this, specifically the correct choice for the ω frequency, especially in the case of virtual orbitals [141]. A damping factor is added to the frequency component, which can be used to turn it into the standard Breit interaction (2.21).

QED self-energy and vacuum polarization. Both of these contributions reduce to single-particle operators that can be toggled on separately. As they form the core of this work, they are discussed in more detail in Section 4.2.

Normal and special mass shifts. GRASP allows including lowest-order nuclear motional corrections [142] in the CI Hamiltonian. Physically, these corrections originate from relaxing the implicit assumption that the nucleus is infinitely heavy and therefore can not move away from the origin. Instead, a nucleus actually has a finite mass M, which can be taken into account with the following operators

$$H_{\mathrm{NMS}} = \frac{c\boldsymbol{\alpha}\cdot\boldsymbol{p} + c^2(\beta-1)}{M}, \quad H_{\mathrm{SMS}} = \frac{\boldsymbol{p}_1\cdot\boldsymbol{p}_2}{M}$$

which are referred to as the *normal mass shift* (NMS) and *special mass shift* (SMS), respectively. They are also one-particle and two-particle operators, respectively.

Including Breit in the variational part of GRASP has been long discussed in the GRASP community, and there are other MCDHF programs that do that [93]. The challenge is largely technical — the MCDHF part of GRASP would require significant refactoring for this to happen. The same applies to the other corrections as well, even though expanding the MCDHF part with a single-particle operator like a QED correction would likely be easier.

4.2 QED corrections in the GRASP software

This section deals with the mathematical details of the QED implementation into the GRASP program package, including the general framework of incorporating new operators into the CI code. It aims to aid in understanding the code implementation, discussed in more detail in Section 4.4.

Theoretical framework

The implementation of the various terms of the Hamiltonian in the CI step of GRASP reduces to computing many-particle matrix elements given the single- or two-particle matrix elements of the operators used. Both the effective self-energy operators used in this thesis and the vacuum polarization potentials are single-particle operators which can be represented (in the single particle subspace spanned by $|n\kappa m\rangle$; see Eq. (4.1)) as

$$A = \sum_{\substack{n,\kappa,m \\ n',\kappa',m'}} A(n,\kappa,m;n',\kappa',m') \, |n\kappa m\rangle \, \langle n'\kappa' m'| \tag{4.3}$$

where only the matrix elements $A(n,\kappa,m;n',\kappa',m')$ need to be known to include such an operator in the calculation of the full many-particle matrix elements of the Hamiltonian.

If it is further assumed that the operator A is invariant under total angular rotation generated by $\boldsymbol{J}$, and hence forms an irreducible tensor of rank 0, then the Wigner-Eckart theorem allows the number of matrix elements to be reduced further. An irreducible tensor can be written in terms of m-independent reduced matrix elements $\langle n\kappa||A||n'\kappa'\rangle$ and the Clebsch-Gordan coefficients[8] as

$$\langle n\kappa m|A|n'\kappa' m'\rangle = C(j,0,j';m,0,m')\delta_{\kappa\kappa'} \, \langle n\kappa||A||n'\kappa'\rangle$$

That, together with the properties of the Clebsch-Gordan coefficients[9], can be used to show that the matrix elements of A are independent of m and diagonal in both m and κ, and so can be written in terms of the reduced elements $A(n,n';\kappa)$ as

$$A(n,\kappa,m;n',\kappa',m') = A(n,n';\kappa)\delta_{\kappa\kappa'}\delta_{mm'}$$

The summation in Equation 4.3 therefore reduces to

$$A = \sum_{n,n',\kappa} A(n,n';\kappa) \sum_{m} |n\kappa m\rangle \, \langle n'\kappa m|$$

The one-particle matrix elements can then be used to calculate the many-particle matrix elements of the operators between two CSFs $|\Psi_a\rangle$ and $|\Psi_b\rangle$

$$\langle \Psi_a|A|\Psi_b\rangle = \sum_{n,n',\kappa} c(a,b;n,n',\kappa)A(n,n',\kappa) \tag{4.4}$$

where the coefficients $c(a,b;n,n',\kappa)$ are determined by the occupations and couplings of the CSFs. The routines for calculating these coefficients have long existed in GRASP and hence the rather complicated details of how such coefficients are calculated [137, 138, 139] are not relevant here. For the implementation of the QED operators, it is simply sufficient to determine the corresponding single-particle matrix elements and combine them with the values provided by the existing GRASP routines.

First-order perturbation theory

A new feature of the updated GRASP implementation is that if a particular QED term is not added to the CI matrix, its contribution can still be evaluated perturbatively to first-order. After

[8] See e.g. [76] equation (5.1) and (5.5) for the definition of irreducible tensors — setting $L = 0, M = 0$ yields the invariance condition. The Wigner-Eckart theorem is discussed in Chapter 5.19.

[9] Specifically that a Clebsch-Gordan coefficient $C(j_1,j_2,j;m_1,m_2,m)$ is zero unless $|j_1 - j_2| \leq j \leq j_1 + j_2$ and $m_1 + m_2 = m$.

the many-electron wavefunction for each state $|n\rangle$ has been determined, given as the expansion

$$|n\rangle = \sum_a C_a^n \, |\Psi_a\rangle$$

in terms of the coefficients C_a^n, the perturbative correction can be evaluated as an expectation value of the corresponding operator

$$\langle n|\hat{A}|n\rangle = \sum_{a,b} C_a^{n*} C_b^n \, \langle \Psi_a|A|\Psi_b\rangle$$

with the help of the operator's many-particle matrix elements $\langle \Psi_a|A|\Psi_b\rangle$.

Effective operators for QED self-energy

Unlike for the vacuum polarization, there is no immediately obvious way to convert QED perturbative self-energy diagrams into an effective operator that could be incorporated as a contribution in the Hamiltonian into multi-electron calculations. In the 2018 version of GRASP, the self-energy contribution is incorporated as a perturbative shift, using values tabulated for the $F(Z\alpha)$ function [68], related to the energy shift of hydrogenic orbitals as

$$\Delta E_{\mathrm{SE}}^{\mathrm{H}} = \frac{\alpha mc^2}{\pi} \frac{(Z\alpha)^4}{n^3} F(Z\alpha)$$

and which have been calculated to high accuracy for the hydrogenic orbitals from first principles in perturbative QED [71, 143]. These hydrogenic values for each orbital $E_{\mathrm{SE}}^{\mathrm{H}}$ are then scaled[10] to obtain the energy shifts E_{SE} of GRASP orbitals according to

$$\Delta E_{\mathrm{SE}} = \frac{\langle \psi|\psi\rangle_{r<R_0}}{\langle \psi^{\mathrm{H}}|\psi^{\mathrm{H}}\rangle_{r<R_0}} \Delta E_{\mathrm{SE}}^{\mathrm{H}}$$

where $|\psi\rangle$ is the GRASP orbital, $|\psi^{\mathrm{H}}\rangle$ the corresponding hydrogenic orbital, and the notation $\ldots_{r<R_0}$ implies that the overlap integral only runs up to $R_0 = 0.0219\ a_0$. Those contributions for each of the orbitals are simply then added up based on the occupation number of the orbital in the given many-electron state.

As part of this work, new effective self-energy operators were incorporated into the CI program of GRASP. The operators are all treated the same way in terms of being represented by their single-particle matrix elements for a given set of GRASP orbitals. Adding a new operator is simply a matter of figuring out how to evaluate the matrix elements for the GRASP orbitals represented on the underlying radial grid.

Self-energy correction from hydrogenic shifts

This is a modified version of the self-energy contribution described above, existing in GRASP since its very early days, originally as a perturbation to the MCDHF solutions. Newer GRASP versions have included an additional option that allows for the energy shifts to be included in the CI matrix, where they are added to the diagonal elements of the CI matrix. The version developed in this work effectively replicates that behaviour, where the energy shift for each

[10] Various ways of scaling the hydrogenic self-energy values have been reviewed in the 2013 paper by Lowe, Chantler, and Grant [73]. Equation (A.3) is the method used in the 2018 version of GRASP.

orbital $E^{\mathrm{SE}}_{n\kappa}$ is simply used to populate the diagonal of the one-particle matrix of the self-energy contribution

$$H^{\mathrm{SE}}_{n\kappa,n'\kappa'} = E^{\mathrm{SE}}_{n\kappa}\delta_{nn'}\delta_{\kappa\kappa'} \tag{4.5}$$

and the off-diagonal elements are kept zero. If the self-energy contribution is evaluated in first-order perturbation theory, it corresponds to the old standard implementation where it is evaluated perturbatively.

Gaussian potential due to Pyykkö and Zhao

In 2003, Pyykkö and Zhao [18] proposed an effective local Gaussian potential for the self-energy of the form

$$\Phi_{\mathrm{SE}}(r) = B(Z)\exp\left[-\beta(Z)\ r^2\right]$$

where B and β are fitted to quadratic polynomials $B(Z) = b_0 + b_1 Z + b_2 Z^2$ and $\beta(Z) = \beta_0 + \beta_1 Z + \beta_2 Z^2$. The coefficients b_n and β_n were fitted such that the potential would reproduce the shifts due to self-energy in hydrogenic and lithium-like systems that had been determined by more accurate calculations in earlier works. The numerical values of the coefficients are given in Table 2 of their article. The evaluation of the matrix elements of the operator reduces to an integral over the potential

$$H^{\mathrm{SE}}_{n\kappa,n'\kappa'} = \delta_{\kappa\kappa'}\int_0^\infty \mathrm{d}r\ \Phi_{\mathrm{SE}}(r)\left[P^*_{n\kappa}(r)P_{n'\kappa'}(r) + Q^*_{n\kappa}(r)Q_{n'\kappa'}(r)\right]$$

Potential due to Flambaum and Ginges

A more sophisticated approach to an effective self-energy potential was propsed in 2005 by Flambaum and Ginges [19]. While this also involves a fit to self-energy energy shifts of hydrogen-like systems tabulated in the literature from perturbative calculations, parts of the operator are derived more directly from QED [144]. The potential is split into three parts,

$$\Phi_{\mathrm{SE}}(r) = \Phi_g(r) + \Phi_f(r) + \Phi_l(r)$$

where the Φ_g, Φ_f and Φ_l terms are referred to as the magnetic form factor, the electric form factor and the low energy contribution, respectively. The magnetic and electric form factor terms are dependent on the nuclear potential $V(r)$ and are given by

$$\Phi_g(r) = \frac{i\alpha}{4\pi m}\boldsymbol{\gamma}\cdot\nabla\left[V(r)\left(\int_1^\infty \frac{e^{-2trm}}{t^2\sqrt{t^2-1}}\,\mathrm{d}t - 1\right)\right] \tag{4.6}$$

$$\Phi_f(r) = -A(Z,r)V(r)\ \frac{\alpha}{\pi}\int_1^\infty \frac{1}{\sqrt{t^2-1}} \times\left\{\left(1-\frac{1}{2t^2}\right)\left[\ln(t^2-1) + 4\ln\left(\frac{1}{Z\alpha}+\frac{1}{2}\right)\right] - \frac{3}{2} + \frac{1}{t^2}\right\}e^{-2trm}\,\mathrm{d}t$$

for the electric form factor contribution, where

$$A(Z,r) = (1.071 - 1.976x^2 - 2.128x^3 + 0.169x^4)\frac{mr}{mr + 0.07Z^2\alpha^2}$$

and $x = (Z - 80)\alpha$. Finally, the low energy contribution is given by

$$\Phi_l(r) = -\frac{B(Z)}{e} Z^4 \alpha^5 mc^2 e^{-Zr/a_B}$$

where $B(Z) = 0.074 + 0.35Z\alpha$. The coefficients of $A(Z, r)$ and $B(Z)$ were determined by fitting the self-energy shifts to earlier calculations of radiative shifts for hydrogenic atoms to improve the results for the high principal quantum number energy levels [72, 125].

The evaluation of the matrix elements for the electric form factor and low energy contribution is the same as for the Pyykkö & Zhao potential, as it is a simple local radial potential

$$H^{\mathrm{SE},i}_{n\kappa,n'\kappa'} = \delta_{\kappa\kappa'} \int_0^\infty \mathrm{d}r\ \Phi_i(r) \left[P^*_{n\kappa}(r) P_{n'\kappa'}(r) + Q^*_{n\kappa}(r) Q_{n'\kappa'}(r) \right], \quad i = f, l$$

The integration of the magnetic form factor is slightly different. By combining everything under the gradient in (4.6) into a single function

$$f(r) = \frac{\alpha}{4\pi m} V(r) \left(\int_1^\infty \frac{e^{-2trm}}{t^2 \sqrt{t^2 - 1}}\, \mathrm{d}t - 1 \right)$$

and noting that $\nabla f(r) = f'(r)\hat{\boldsymbol{r}}$, the magnetic form factor can be written as

$$\Psi_g(\boldsymbol{x}) = i\boldsymbol{\gamma} \cdot \nabla f(r) = i\boldsymbol{\gamma} \cdot \hat{\boldsymbol{r}} f'(r)$$

where it becomes more clear that this is an *odd operator* [62, Section 6.2]. The integral for the matrix element becomes

$$H^{\mathrm{SE},g}_{n\kappa,n'\kappa'} = \delta_{\kappa\kappa'} \int_0^\infty \mathrm{d}r\ f'(r) \Big[P^*_{n\kappa}(r) \cdot Q_{n'\kappa'}(r) + Q^*_{n\kappa}(r) \cdot P_{n'\kappa'}(r) \Big]$$

where instead of the f it is necessary to evaluate its derivative

$$f'(r) = \frac{\alpha}{4\pi m} \left[V'(r) \left(\int_1^\infty \frac{e^{-2trm}}{t^2 \sqrt{t^2 - 1}}\ \mathrm{d}t - 1 \right) - 2mV(r) \int_1^\infty \frac{e^{-2trm}}{t\sqrt{t^2 - 1}}\ \mathrm{d}t \right]$$

instead. It also needs the derivative of the nuclear potential $V'(r)$ to be available, and the inner integrals have to be evaluated for each r grid point. The total matrix element for this potential becomes then simply a matter of summing up the matrix elements for the individual contributions

$$H^{\mathrm{SE}}_{n\kappa,n'\kappa'} = H^{\mathrm{SE},g}_{n\kappa,n'\kappa'} + H^{\mathrm{SE},f}_{n\kappa,n'\kappa'} + H^{\mathrm{SE},l}_{n\kappa,n'\kappa'}$$

Self-energy model operator due to Shabaev and Yerokhin

In 2013, Shabaev et.al. [20] proposed the following non-local operator based on the spectral representation of the SE operator for modelling the self-energy contribution

$$\hat{h}_{\mathrm{SE}} = \sum_\kappa V_\kappa(r) \hat{P}_\kappa + \sum_{i,j} |\phi_i\rangle\, B_{ij}\, \langle\phi_j|$$

where $V_\kappa(r) = A_\kappa \exp\left(-rmc/\hbar\right)$, $\hat{P}_\kappa$ is the projection operator onto the κ angular momentum subspace, and the states $|\phi_i\rangle$ are a particular model basis proposed in the paper. The constants

A_κ are determined by requiring the $V_\kappa(r)$ to reproduce the self-energy shift of the lowest hydrogenic state of the specified angular momentum. B_{ij} are calculated from the exact hydrogenic overlaps. This local operator significantly improves the results over a pure spectral representation explored by Dyall with little success as the chosen orbitals did not span a large enough space [145].

While more complex than a simple local potential, this operator can quite easily be included in many-electron calculations. The authors have published a set of Fortran routines under the name QEDMOD [146, 147] which contain all the necessary code to evaluate this operator. Specifically, the code allows the large and small components $P_{n\kappa}(r)$ and $Q_{n\kappa}(r)$ to be transferred onto its own numerical grid and to evaluate the effect of the operator on those orbitals. The output is another large-small component pair $P'_{n\kappa}(r)$ and $Q'_{n\kappa}(r)$, formally defined as

$$\frac{1}{r}\begin{pmatrix} P'_{n\kappa}(r)\chi_{\kappa m}(\Omega) \\ -iQ'_{n\kappa}(r)\chi_{-\kappa m}(\Omega)\end{pmatrix} = \hat{h}_{\mathrm{SE}}\left[\frac{1}{r}\begin{pmatrix} P_{n\kappa}(r)\chi_{\kappa m}(\Omega) \\ -iQ_{n\kappa}(r)\chi_{-\kappa m}(\Omega)\end{pmatrix}\right]$$

The authors do use a different sign convention for the small component compared to the GRASP convention of Equation (4.1), which needs to be taken into account when transferring GRASP orbitals into QEDMOD and back. To evaluate the single-particle matrix elements, it is simply necessary to evaluate the overlap integral between the left-hand state and the transformed right-hand state

$$H^{\mathrm{SE}}_{n\kappa,n'\kappa'} = \langle n\kappa|\hat{h}_{\mathrm{SE}}|n'\kappa'\rangle = \delta_{\kappa\kappa'}\int_0^\infty \mathrm{d}r\,\left[P^*_{n\kappa}(r)P'_{n'\kappa'}(r) + Q^*_{n\kappa}(r)Q'_{n'\kappa'}(r)\right]$$

Vacuum polarization potentials

While the vacuum polarization was not updated in terms of adding any higher-order contributions, it was significantly refactored. Hence, it is useful to document here the underlying mathematical expressions. As discussed in Section 2.5, vacuum polarization reduces to a potential $V_{\mathrm{vp}}(\boldsymbol{x})$, which if the nuclear charge distribution is radially symmetric, also becomes a radially symmetric function (and hence a rank-0 tensor). The evaluation of the matrix elements becomes a simple radial integral over the potential with the orbitals

$$H^{\mathrm{VP}}_{n\kappa,n'\kappa'} = \delta_{\kappa\kappa'}\int_0^\infty \mathrm{d}r\; V_{\mathrm{vp}}(r)\left[P^*_{n\kappa}(r)P_{n'\kappa'}(r) + Q^*_{n\kappa}(r)Q_{n'\kappa'}(r)\right]$$

Uehling contribution

Vacuum polarization in first-order in both α and $Z\alpha$ was originally derived by Uehling in 1935 [65], and actually predates QED. A convenient explicit expression for the Uehling potential given an arbitrary nuclear charge distribution $\rho(\boldsymbol{x})$ [148] is

$$V_{\mathrm{U}}(\boldsymbol{x}) = -\frac{2}{3}\frac{Z\alpha e^2}{\pi}\int \mathrm{d}^3\boldsymbol{x}'\,\frac{\rho(\boldsymbol{x}')}{|\boldsymbol{x}-\boldsymbol{x}'|}K_1\left(\frac{2}{\lambda\!\!\!^{-}_{\mathrm{e}}}|\boldsymbol{x}-\boldsymbol{x}'|\right) \tag{4.7}$$

where $\lambda\!\!\!^{-}_{\mathrm{e}} = \frac{\hbar}{m_{\mathrm{e}}c} = \alpha\cdot a_0$ is the reduced electron Compton wavelength and $K_1(x)$ is defined as

$$K_1(x) = \int_1^\infty \mathrm{d}t\; e^{-xt}\left(\frac{1}{t^2}+\frac{1}{2t^4}\right)\sqrt{t^2-1}$$

In the case of a point nucleus, the potential (4.7) the outer integral disappears and the potential becomes

$$V_{\mathrm{U}}(\boldsymbol{x}) = -\frac{2}{3}\frac{Z\alpha e^2}{\pi|\boldsymbol{x}|}K_1\left(\frac{2|\boldsymbol{x}|}{\lambda\!\!\!^{-}_{\mathrm{e}}}\right)$$

Källén-Sabry contribution

Källén and Sabry considered the contribution of vacuum polarization in the order α^2, but first order in $Z\alpha$. [60] The derivation of the potential form of the contribution was first done in [149] for a point nuclear charge. The expressions used in GRASP, which have been generalized to arbitrary spherically symmetric charge distributions [148], are

$$V_{\mathrm{KS}}(\boldsymbol{x}) = -\frac{e^2\alpha^2}{\pi^2}\int \mathrm{d}^3\boldsymbol{x}'\frac{\rho(\boldsymbol{x})}{|\boldsymbol{x}-\boldsymbol{x}'|}L_1\left(\frac{2}{\lambda_e}|\boldsymbol{x}-\boldsymbol{x}'|\right)$$

where L_1 is

$$\begin{aligned}L_1(x) = \int_1^\infty \mathrm{d}t\ e^{-xt}\Bigg[&\left(-\frac{13}{54t^2}-\frac{7}{108t^4}-\frac{2}{9t^6}\right)\sqrt{t^2-1}\\ &+\left(\frac{44}{9t}-\frac{2}{3t^3}-\frac{5}{4t^5}-\frac{2}{9t^7}\right)\ln(t+\sqrt{t^2-1})+\left(-\frac{4}{3t^2}-\frac{2}{3t^4}\right)\sqrt{t^2-1}\ln(8t^3-8t)\\ &+\left(\frac{8}{3t}-\frac{2}{3t^5}\right)\int_t^\infty \mathrm{d}y\left(\frac{3y^2-1}{y(y^2-1)}\ln(y+\sqrt{y^2-1})-\frac{\ln(8y^3-8y)}{\sqrt{y^2-1}}\right)\Bigg]\end{aligned}$$

Vacuum polarization in higher orders of $Z\alpha$

The evaluation of the $(Z\alpha)^{3+}$ terms for the first order vacuum polarization (i.e. full $(Z\alpha)$ contribution on the Uehling level) was done by Wichmann and Kroll in the 1950s [66]. The equations 3–5 in [150] put their work into a form that is relatively straightforward to evaluate. However, as of now, this contribution has not been incorporated into GRASP yet.

4.3 Implementation benchmark: Be-like system

To verify that the refactored `rci` program of GRASP and the new QED operators work correctly and produce both numerically and physically reasonable results, the updated code was systematically tested on the Be-like isoelectronic sequence (i.e. ions with four electrons). This particular system is convenient for the following reasons:

1. The system is simple, as it only has four electrons, and so it is computationally easy to include all the correlation in a full CI calculations.

2. The ground state of the system ($1s^2 2s^2$) has a closed $2s$ shell (as opposed to the 3-electron Li-like isoelectronic sequence), leading to a $J = 0$ ground state.

3. Detailed results with a slightly older version of GRASP were reported by Fischer et.al. for a few specific nuclear charges [127], making it easy to at least qualitatively compare with known reference values.

Essentially, the Be-like isoelectronic sequence should be complicated enough in terms of correlation that any issues that are due to the integration of the QED single-particle operators into a MCDHF procedure would show up. At the same time, it is simple enough computationally that there is no need to worry about the computational cost.

Computational procedure

The goal was to compute the Be-like isoelectronic sequence for $4 \leq Z \leq 118$. Compared to the reference values by Fischer et.al., where only three different elements were computed and U^{88+} ($Z = 92$) being the heaviest, the benchmark involves a full scan in Z, all the way up to oganesson with $Z = 118$ — the heaviest experimentally discovered nucleus [151].

The ground state for most Be-like ions is $1s^2 2s^2\ {}^1S_0$ and it would, therefore, make sense to use the single $J = 0$ CSF generated from the $1s^2 2s^2$ configuration as the multireference space. However, it turns out that, at higher nuclear charges, this no longer yields a good description of the ground state, and so most of the calculations were performed with a multireference space that contains all substitutions from $1s^2 2s^2$ into $2p$ (a total of 8 CSFs).

When optimizing the orbitals, GRASP is only able to variationally optimize the orbitals for the Dirac-Coulomb Hamiltonian (i.e. Breit or QED operators are not included in this step). Following the standard GRASP procedure, the correlation orbitals are generated in a layer-by-layer fashion, where after the $n = 1$ and $n = 2$ multireference orbitals have been optimized, the rest are generated and optimized one n-shell at a time, with orbitals with lower n values frozen. Each new layer of orbitals will be orthogonal to all the existing, frozen orbitals. The CSFs for electron correlation are generated by performing a maximum number of substitutions (4) from the multireference configuration into the whole correlation space, generating a complete basis for a given angular momentum and parity (up to a set of orbitals). No restrictions are applied on the angular momentum of the orbitals.

In the following discussion and figures, the configuration state list (CSL) used for a calculation will be referred to by the `multiref` and `cas-nN` shorthands. The former refers to the relevant multireference space, whereas `cas-nN` refers to the spaces that include all the substitutions into the correlations orbitals with principal quantum number $n \leq$ `N` (i.e. the full CI *complete active space*, with orbitals up to $n =$ `N`).

QED corrections

The contributions that go beyond the Dirac-Coulomb Hamiltonian (including Breit and QED) were calculated with the CI program `rci` that had been modified as part of this work. While this can be thought of as an extra CI step, the calculations use the same space as the variational calculations (as the space is already complete), and therefore this step does not actually include any additional correlation. Conceptually, these calculations are a simple re-diagonalization of the full many-particle Hamiltonian matrix with some or all of the terms in the Hamiltonian

$$H = H_{\mathrm{D}} + H_{\mathrm{C}} + H_{\mathrm{B}} + H_{\mathrm{VP}} + H_{\mathrm{SE}} + H_{\mathrm{MS}}$$

enabled. The different terms correspond to (1) free Dirac + nuclear potential, (2) 2-particle Coulomb interaction, (3) 2-particle Breit interaction, (4) QED vacuum polarization (Uehling and Källén-Sabry), (5) effective QED self-energy operators, and (6) normal and special mass shifts, respectively. To see the effect of each term and to also compare with reference values reported by Fischer et.al. [127], the calculation is run multiple times with different terms enabled in each case. In particular, the following Hamiltonians were used

1. $H_{\mathrm{D}} + H_{\mathrm{C}}$
2. $H_{\mathrm{D}} + H_{\mathrm{C}} + H_{\mathrm{B}}$
3. $H_{\mathrm{D}} + H_{\mathrm{C}} + H_{\mathrm{B}} + H_{\mathrm{VP}}$
4. $H_{\mathrm{D}} + H_{\mathrm{C}} + H_{\mathrm{B}} + H_{\mathrm{VP}} + H_{\mathrm{SE}}$

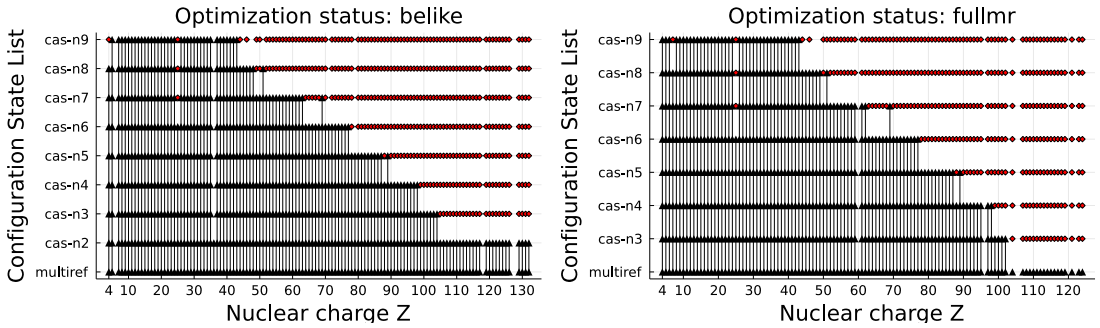

Figure 4.1: Illustration of which CSL and Z combinations successfully managed to optimize the orbitals with RMCHDF, for both the $1s^2 2s^2$ (left) and the full $n = 2$ (right) multireference spaces. The black triangles indicate successful optimisations. The red diamonds are the cases where RMCDHF apparently succeeded but produced nonsensical results and so were removed by hand. Missing points indicate cases where RMCDHF failed with an error.

A contribution from the extra term is defined as the difference between the calculation with and without the extra term. In each case, all the omitted terms can also be evaluated using first order perturbation theory. The mass shift terms were not evaluated as the routines for those were unchanged.

Nuclear models

The atomic nuclear charge distributions were modelled with the standard two-parameter Fermi model, which takes into account the finite extent of the physical nucleus extremely well [152]. The exact parameters, which get stored in the `isodata` files, are determined with GRASP's `rnucleus` program using a reference RMS value and the fixed value of $t = 2.3$ fm for one of the parameters of the Fermi model. The RMS values are mostly sourced from the tables by Angeli and Marinova [84], but for the nuclei that do not have tabulated values, including for all $Z \geq 97$, GRASP falls back to the empirical formula (2.34) derived by Johnson and Soff [85].

Manual data cleanup

The `rmchdf` program in GRASP regularly has problems converging the orbitals, either by exiting with an error code or more often silently failing and producing nonsensical results (e.g. containing `NaN`s). Because of the latter, the RMCHDF optimisation results were further cleaned by hand by removing the problematic CSL and Z combinations. No significant effort was made to fix these particular optimisations, since there are simply too many cases that result in an error. It does, however, indicate that GRASP would greatly benefit from work addressing the reliability of the optimisation routines in the `rmcdhf` program.

The result of the filtering is illustrated in Figure 4.1. Generally, at higher Z values, it becomes more and more difficult to converge orbitals with a higher n quantum number. This is consistent with the results presented in Table 8 in the reference paper by Fischer et.al. [127], where it can be seen that the maximum n attained is 8, 6 and 5, for $Z = 20, 54, 92$, respectively.

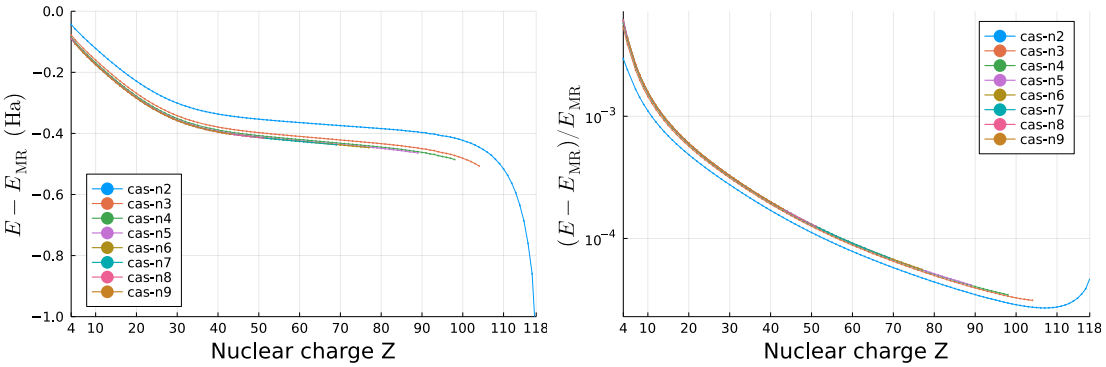

Figure 4.2: Correlation energy of the Be-like isoelectronic sequence assuming a single-determinant $1s^2 2s^2$ multireference space at the Dirac-Coulomb level for as a functions of nuclear charge Z. The figure on the left shows it on an absolute scale (in Hartrees), whereas the figure on the right shows it relative to the multireference energy. The different colors indicate the number of correlation orbitals included.

Correlation in Be-like systems

After some preliminary calculations, it became clear that the correlation behavior of the Be-like isoelectronic sequence at high Z values $Z \geq 100$ is more complicated than anticipated. For that reason, before going into the QED benchmarks themselves, it is necessary to give a detailed description of the issues and workarounds regarding electron correlation. This is important for understanding certain artefacts that appear in the QED corrections at those high nuclear charges.

The problems caused by the correlation behaviour are illustrated in Figure 4.2, which shows the correlation energy as a function of Z when using $1s^2 2s^2$ for the multireference space. Starting from the low Z values, we see that, relative to the multireference energy, the correlation energy decreases with increasing nuclear charge. This is intuitively expected, as with a deepening nuclear potential well, the orbital splittings become larger, and hence electron correlation becomes less important. However, what is also clearly visible is that at high nuclear charges the correlation energy suddenly increases. This, in turn, will introduce artefacts into the Z-behaviour of the QED corrections, and therefore it is important to understand how to mitigate this.

It is clear that most of the correlation comes from including the additional $n = 2$ orbitals, which is seen from the fact that the correlation values in Figure 4.3 are about an order of magnitude lower than in Figure 4.2. However, as it was impossible to converge any correlation orbitals at the higher nuclear charges, it is uncertain from that how this affects the bend in the correlation energy. On the other hand, Figure 4.4 does show that the correlation contribution from the $n = 2$ orbitals significantly reduces the ground state energy in the $Z \geq 110$ region.

It turns out that the reason for this correlation effect is that the ground state configuration changes at very high Z. This is illustrated by Figure 4.5, which shows the mixing probabilities ($|c_i|^2$ of the mixing coefficients) of the dominant CSFs in the ground state. Slightly above $Z = 120$ the ground state configuration changes from $1s^2_{1/2} 2s^2_{1/2}$ to $1s^2_{1/2} 2p^2_{1/2}$. However, already earlier the $1s^2_{1/2} 2s^2_{1/2}$ configuration starts mixing in more strongly, creating a noticeable shift in the correlation energy at a lower Z value.

This change in the ground state configuration is related to the energy of the $2p_{1/2}$ orbital falling below the energy of the $2s_{1/2}$ orbital. This can be seen in Figure 4.6, which shows the orbital energies optimized for the ground state with a full $n = 2$ multireference, i.e. with the

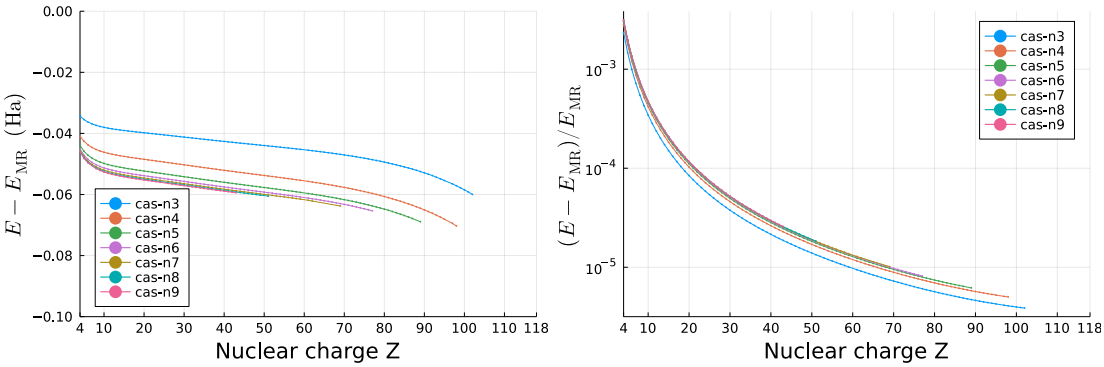

Figure 4.3: Same as Figure 4.2, but using the all $n = 2$ configurations multireference. This leads to a correlation contribution that is about an order of magnitude smaller as most of it is taken into account by the multireference space already.

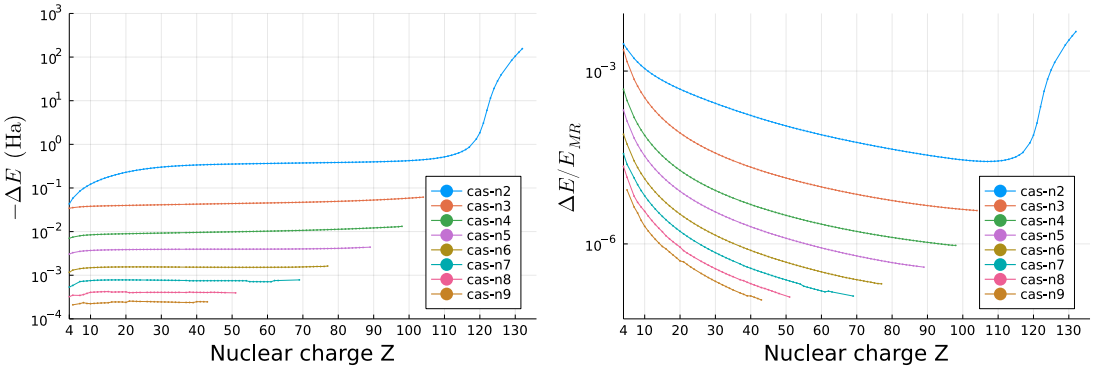

Figure 4.4: Change in Be-like ground state energy by adding in another layer of orbitals for extra correlation. The figure on the left shows the change on an absolute scale, whereas the figure on the right shows the energy relative to the multireference energy.

$2p_{1/2}$ and $2p_{3/2}$ orbitals and related configurations included in the multireference space.

Interestingly, it is the finite extent of the nuclear model that leads to this. In a hydrogenic system for a point nucleus, with QED and other extra effects neglected, the $2s_{1/2}$ and $2p_{1/2}$ are exactly degenerate. The introduction of interactions via the mean-field approach in multi-electron calculations breaks that degeneracy, leading to a slightly higher energy of $2p_{1/2}$ relative to $2s_{1/2}$, illustrated by Figure 4.7. This is consistent with the interpretation where $2p_{1/2}$ experiences more screening of the nuclear potential than $2s_{1/2}$. On the other hand, in the interaction- and therefore correlation-free hydrogenic picture of orbitals, introducing a nucleus with a finite extent leads to a splitting in the opposite direction, with $2p_{1/2}$ lying lower than $2s_{1/2}$. This can be understood in terms of the finite nucleus pushing the orbital energies up, relative to the point nucleus case, due the nuclear potential well not being as deep, with orbitals that are closer to the origin being pushed up more on average.

This is all illustrated in Figure 4.8, showing how at high enough Z value the $2p_{1/2}$ energy starts dropping (black line), just as in the hydrogenic case with a finite nucleus (green line). These two effects are essentially competing with each other, with the multi-electron effect dominating

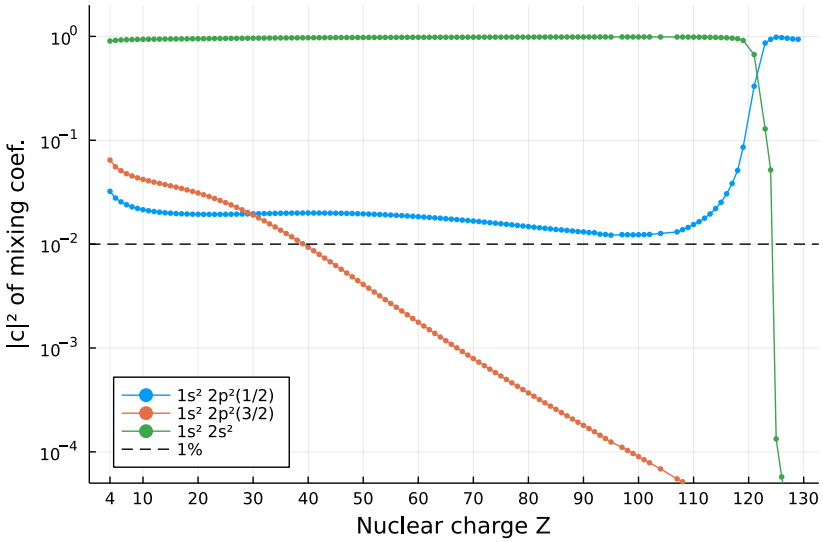

Figure 4.5: Mixing probabilities of the dominant CSFs for the Be-like isoelectronic sequence, obtained from a full $n = 2$ multireference calculation.

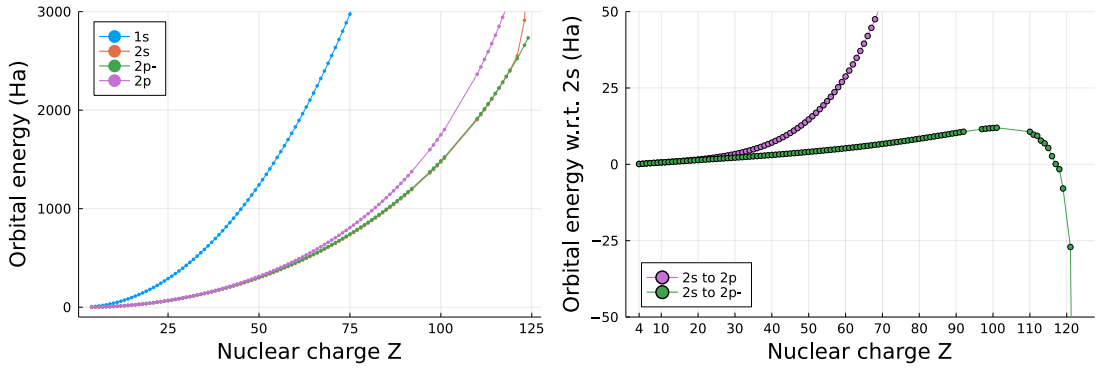

Figure 4.6: Orbital energies for the Be-like isoelectronic sequence, obtained from a multireference calculation with the full $n = 2$ multireference. The figure on the left shows absolute energies, and the figure on the right shows the energies of $2p_{1/2}$ and $2p_{3/2}$ orbitals relative to the $2s_{1/2}$ orbital.

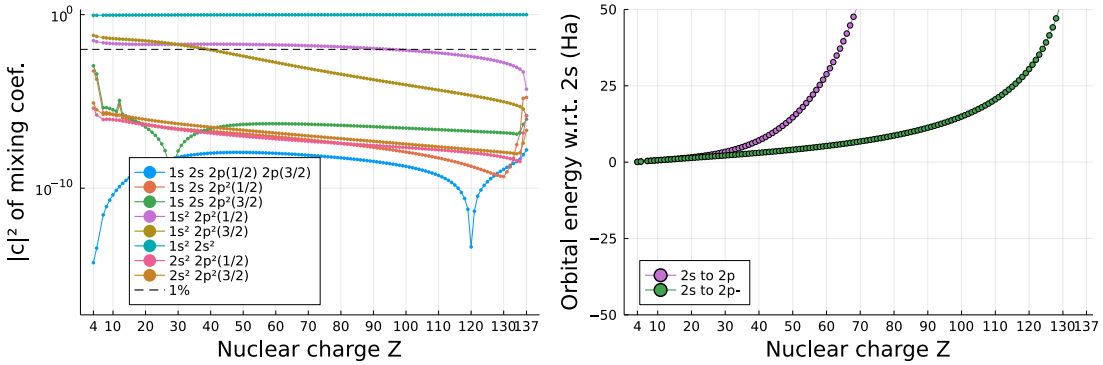

Figure 4.7: CSF mixing probabilities (left) and $2p_{1/2}$ and $2p_{3/2}$ orbital energies (relative to $2s_{1/2}$; right) for a *point nucleus* in the Be-like isoelectronic sequence with a full $n = 2$ multireference space.

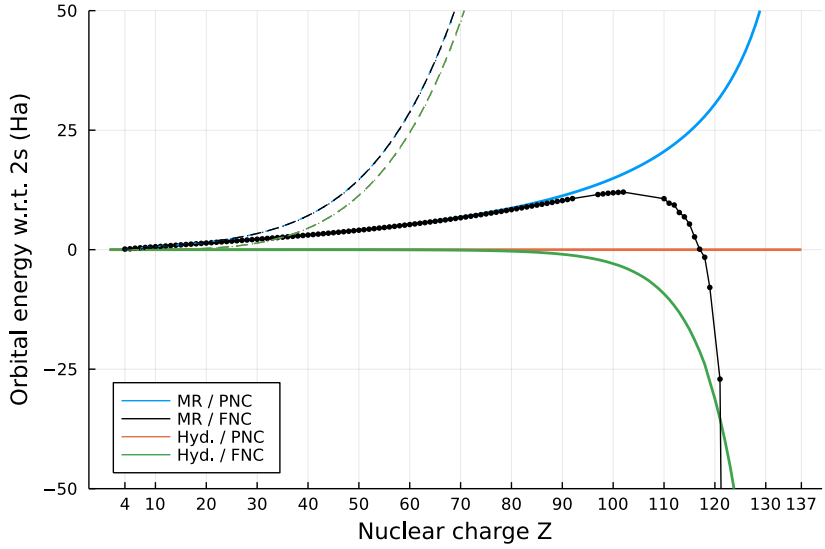

Figure 4.8: Difference in orbital energies for the $2p_{1/2}$ (solid) and $2p_{3/2}$ (dashed) orbitals with respect to $2s_{1/2}$ orbital. The black lines with data points represent the most physical multi-electron and finite nucleus case, showing how the energy of $2p_{1/2}$ eventually falls below $2s_{1/2}$. For comparison, the blue line shows the multi-electron energy but with a point nucleus, the orange line the hydrogenic point nucleus energy (where the two orbitals are degenerate), and the green line shows the hydrogenic energy with a point nucleus (where $2p_{1/2}$ is always lower in energy than $2s_{1/2}$). The hydrogenic energies are for a uniform shell nucleus using the same constant RMS values for all nuclei (the RMS values of $Z = 118$), which is not the same nuclear model as for the multi-electron case (Fermi nucleus with a Z-dependent RMS), but this is fine as the relatively minor difference in the nuclear potential does not lead to any qualitative difference in the energies. The $2p_{3/2}$ energies are shown as dashed lines, where the multi-electron case has a slightly higher energy, but whether it is a finite or a point nucleus does not make a visible difference to the $2p_{3/2}$-$2s_{1/2}$ energy splitting in neither the multi-electron nor the hydrogenic case.

Table 4.1: Dirac-Hartee-Fock energies for the Be-like cases, compared with the reference values by Fischer et.al. [127] for the cases where such values are available. All energies are given in Hartrees.

Z	**This work**	**Fischer 2016**	**Difference**	**Difference (fraction)**
20	-471.959	-471.959	-0.000154836	3.28071e-7
54	-3716.46	-3716.45	0.0084631	-2.27719e-6
92	-12039.7	-12040.1	-0.449982	3.7375e-5

at the lowest Z values that are usually of interest, but the finite nuclear splitting taking over at high nuclear charges. Effectively, a strong enough single-particle Coulomb potential eventually achieves dominance over the many-body mean-field and electron correlation.

Breit and QED contributions to Be-like energies

The focus of this work is on the implementation of additional ways of estimating the effect of the QED self-energy on the energies and states. In order to add new self-energy operators, it was necessary to significantly reorganise the CI program `rci`, where all the additional terms that go beyond the Dirac-Coulomb Hamiltonian are implemented. As GRASP has no unit or integration tests that could automatically check that the reorganisation does not introduce any bugs or errors into the code, it was necessary to verify this manually. This section, therefore, verifies that the existing terms are still yielding the same results as before and also that the new self-energy approaches yield sensical results. This is done by looking at the behaviour of the operators in detail for the Be-like isoelectronic sequence, including keeping an eye on any edge cases where the operators or their implementation may run into problems. To verify the existing contributions (Breit, vacuum polarization and hydrogenic self-energy), the results from the updated program are compared to the results by Fischer et.al. [127].

Dirac-Hartree-Fock DC energies

Table 4.1 shows a small, but consistent difference between the multireference energies calculated with the updated program and the reference values by Fischer et.al. [127]. It should be noted that for $Z = 54$ the sign of the difference is opposite, and it is not immediately obvious where these differences come from. It is always possible that this is simply due to a different convergence criterium. Or, alternatively, it could be due to differences in the exact parametrisation of the nuclear models, which may be likely as the reference data for nuclear models was updated between the GRASP2k version (used for the reference values) and the 2018 GRASP version (used in this work). Most likely, though, it is a combination of multiple factors contributing to these discrepancies.

Electron correlation energy

Figure 4.9 shows the electron correlation energy for the Be-like isoelectronic sequence, depending on the number of correlation orbital shells included in the calculation. The curve is smooth and agrees very well with the reference values, indicating the orbital optimisation reported in the reference paper has been successfully replicated. The bend at very high Z is related to change in the ground state, as discussed in the previous subsection.

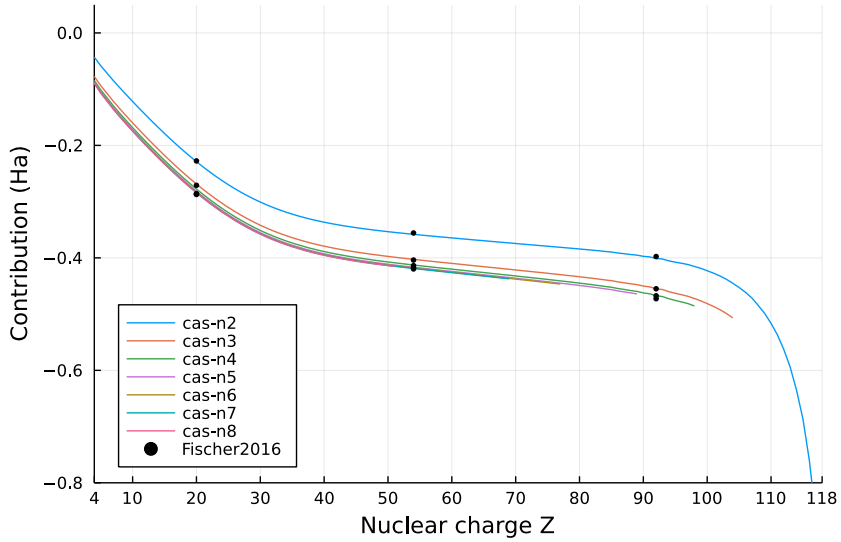

Figure 4.9: Correlation contribution at the DC level for the Be-like isoelectronic sequence with different number of correlation orbital shells. The black dots correspond to the values by Fischer et.al. [127].

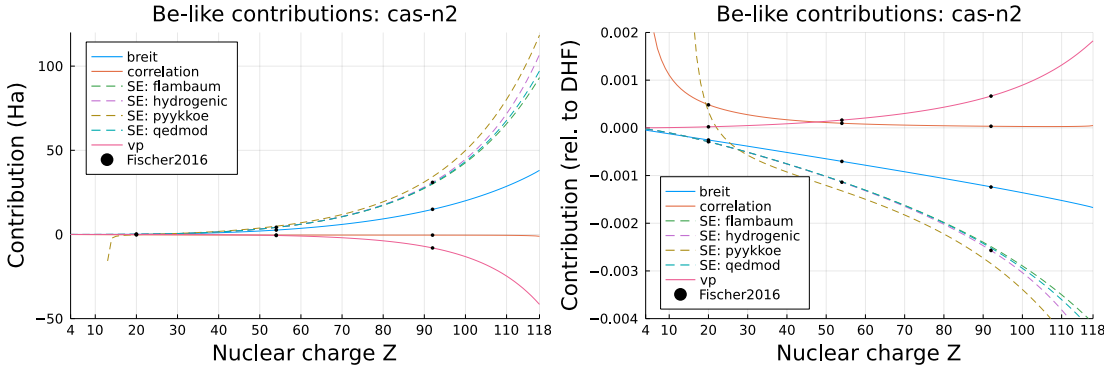

Figure 4.10: Non-variational contributions from `rci` at the $n = 2$ correlation level. The figure on the left shows absolute energies, whereas the value on the right figure are relative to the DHF DC energy. The black dots correspond to the values by Fischer et.al. [127].

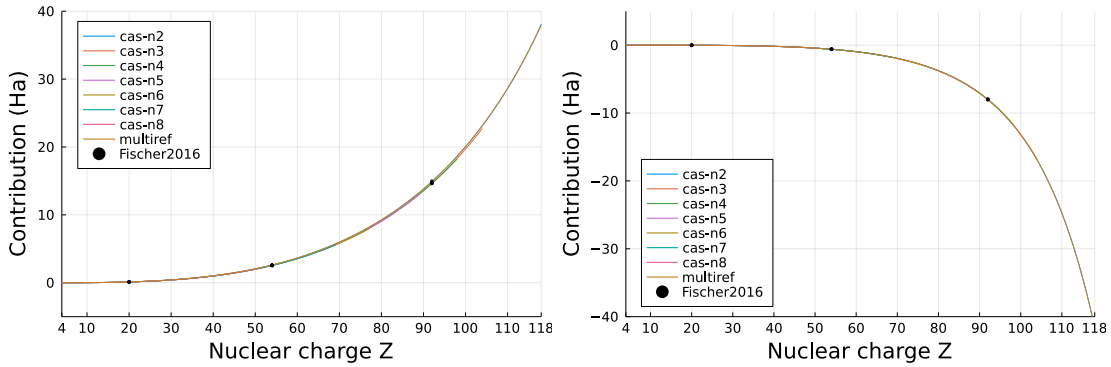

Figure 4.11: Breit (left) and vacuum polarization (right) contributions at various correlation levels in the Be-like isoelectronic sequence. The black dots correspond to the values by Fischer et.al. [127].

Existing contributions

Figure 4.10 shows the electron correlation and all post-Dirac-Coulomb contributions for the Be-like isoelectronic sequence. As the energy scales with Z, it is helpful to look at the energy relative to the MR energy, as is done in that figure on the right. This makes the behaviour and agreement with the reference data more evident also at low Z values, although only at the $n = 2$ CAS correlation level as that is the only case where reference data is available for all three nuclei calculated in the reference article. The figures show a very good agreement with the reference values by Fischer et.al. [127], including for the hydrogenic self-energy. Relatedly, Figure 4.11 shows the Breit and vacuum polarization contributions separately for all correlation levels, demonstrating that the contributions are to quite high precision already determined at multireference level, with very little change due to electron correlation.

New self-energy implementations

As could be seen in Figure 4.10, the different self-energy implementations exhibit noticeable differences at higher Z values. Generally, the Flambaum-Ginges and QEDMOD approaches are consistent and very slightly divergent from the original hydrogenic approach. The simpler Pyykkö-Zhao operator, however, diverges significantly, with especially non-physical behavior at low Z values (including a sign-change). Figure 4.12 shows, as for the Breit and vacuum polarization, the self-energy contributions are also relatively independent of electron correlation and mostly captured already at the multireference level. It also confirms what can be observed in Figure 4.10: the hydrogenic values calculated with the updated code agree very well with the reference values, as one would expect, whereas the new methods show very slight disagreement. The disagreement with the hydrogenic reference values is noticeably worse for the Pyykkö-Zhao operator.

Electron correlation effects

Looking at absolute contributions does not clearly demonstrate the effect of electron correlation on the contributions, as the change due to electron correlation is extremely small relative to the absolute value. Instead, it is helpful to see much the different contributions change when a new layer of correlation orbitals (i.e. all orbitals with a particular principal quantum number n) is

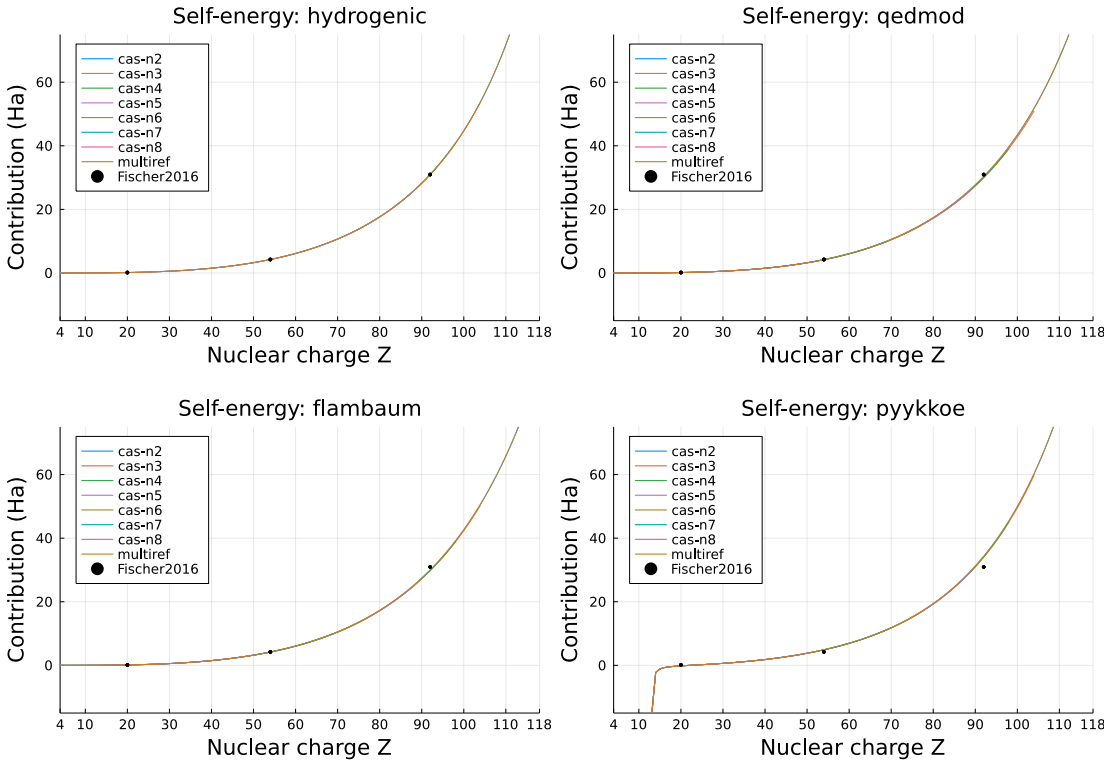

Figure 4.12: QED self-energy contributions for the different implementations at various correlation levels in the Be-like isoelectronic sequence. The black dots correspond to the values by Fischer et.al. [127].

introduced. More precisely, given two different correlation models S_1 and S_2, it is useful to look at the difference

$$\delta(S_2, S_1; H_2, H_1) = \Delta(S_2; H_2, H_1) - \Delta(S_1; H_2, H_1)$$

In the following figures, `cas-n2` refers to the difference between $n = 2$ and the multireference values, whereas all the other ones refer to differences between layers with consecutive n quantum numbers.

The figures will also include the reference values by Fischer et.al. [127]. Something to note about the reference values, though, is that as the values for the multireference and some n-layers were omitted from the table, it is not possible to have reference values for all the cases. Furthermore, in some cases, the reference values for different layers are identical due to the limited number of digits reported in the table in the reference paper, which in turn means that it is not possible to plot those points either as the figures are on a logarithmic scale. For these reasons, the figures also show the uncertainty in the reference values, determined from the number of digits reported in the paper.

Figure 4.13 shows this for the correlation energy, and demonstrated excellent agreement with the reference values even at very high n values. Similarly, the Breit and vacuum polarization contribution (in Figure 4.14) also agree quite well, broadly speaking. However, especially for the vacuum polarization, but also at higher n values for Breit, the reference values are limited by the number of digits, making it unclear how reliable this conclusion is. As a specific example,

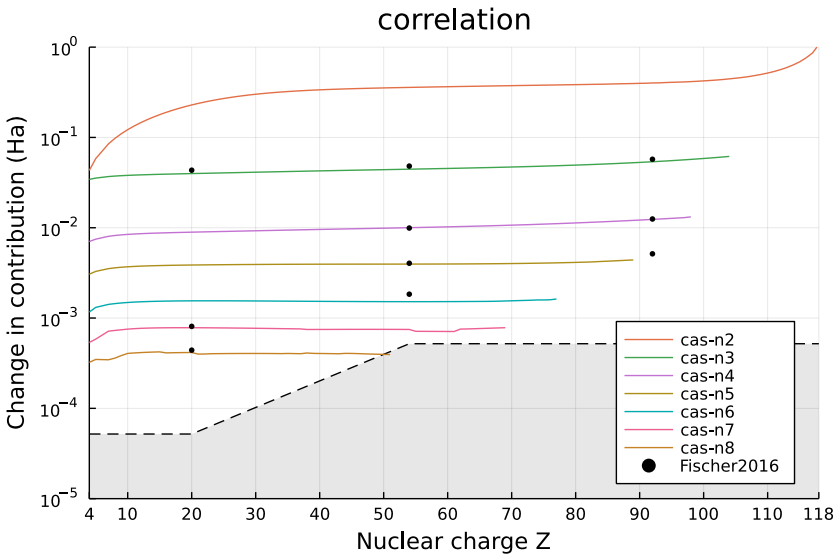

Figure 4.13: Change in the correlation contribution with additional correlation orbital layers of the ground state in the Be-like isoelectronic sequence. The black dots correspond to the values by Fischer et.al. [127]. The grey area bounded by the dashed line indicates the numerical uncertainty of the reference values by Fischer et.al., as determined from the number of digits reported in the paper.

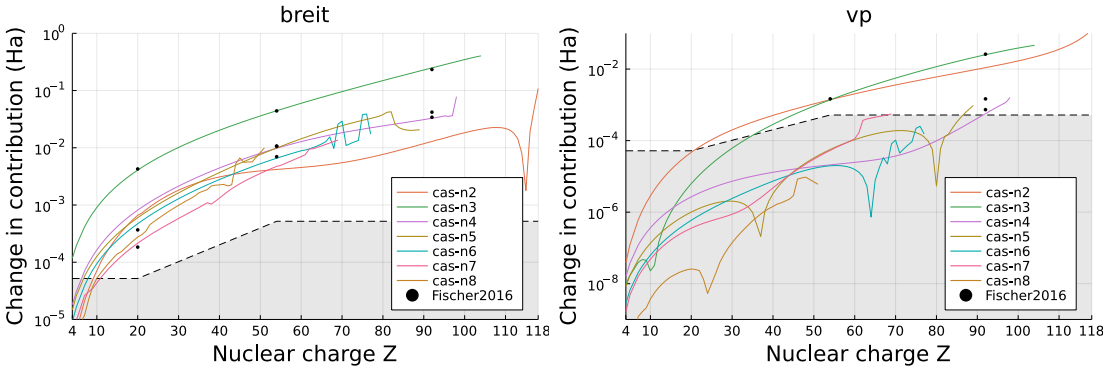

Figure 4.14: Change in the Breit (left) and QED vacuum polarization (right) contributions with additional correlation orbital layers of the ground state in the Be-like isoelectronic sequence. The black dots correspond to the values by Fischer et.al. [127]. The grey area bounded by the dashed line indicates the numerical uncertainty of the reference values by Fischer et.al., as determined from the number of digits reported in the paper.

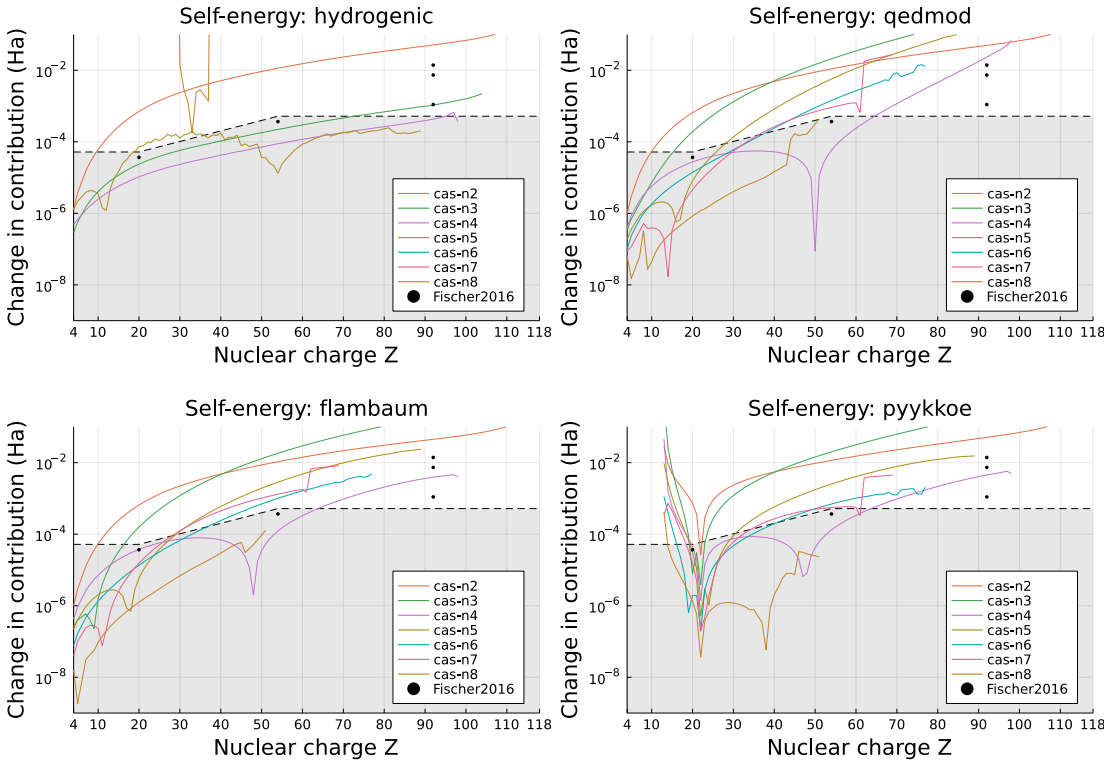

Figure 4.15: Change in the QED self-energy contributions for the different methods with additional correlation orbital layers of the ground state in the Be-like isoelectronic sequence. The black dots correspond to the values by Fischer et.al. [127].

the reference values for $Z = 20$ are completely missing for vacuum polarization, as they are reported to be identical for all correlation layers within the number of digits reported. However, it should be pointed out that this is consistent with the results from this work, with the change in contribution at low Z being far below the number of digits reported in the reference paper.

What can also be seen from the calculations is that at higher n layers the values for the change in contribution become quite noisy, consistent with the optimisation becoming more difficult and unreliable in that regime. However, with this data, it is not possible to say whether this due to optimisation inaccuracies or due to the numerical implementation of the operators introducing that numerical noise at that level of precision, and this should be further investigated.

Finally, Figure 4.15 shows the same information, but for the self-energy contributions. It is clear that the different operators lead to somewhat different results, although the QEDMOD and Flambaum-Ginges operators are qualitatively similar. It is also encouraging to see good agreement in the hydrogenic self-energy case on the `cas-n3` line, even though it is a regime where the uncertainty of the reference values starts becoming significant. The strange behaviour at around $Z = 20$ for Pyykkö-Zhao is due to the sign of the contribution changing, but it should also be pointed out that at high Z values the behaviour of this operator is quite similar to Flambaum-Ginges and QEDMOD.

One thing that should be noted is that the two reference points around 10^{-2} Ha, are actually for layers with larger n values, i.e. the change in the reference self-energy is bigger at higher n.

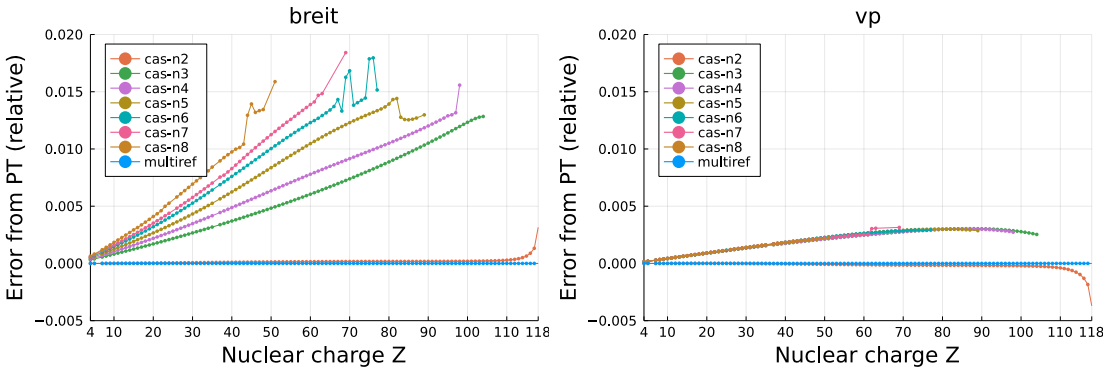

Figure 4.16: Difference in first-order perturbation theory results from the exact contributions, as a fraction of the latter, for Breit (left) and vacuum polarization (right) at different electron correlation levels. The zeroth order Hamiltonian here is DC and DCB, for the Breit and vaccuum polarization, respectively.

This indicates that the hydrogenic self-energy implementation in GRASP2k was not very reliable for high n orbitals. But this is also true for the refactored implementation in GRASP, as the new calculations with values for $n \geq 6$ are highly problematic, indicating that it is very much necessary to restrict the hydrogenic self-energy to only low-n orbitals.

First-order perturbation theory

Finally, as the `rci-qed.pt` program implements first-order perturbation theory estimates for the Hamiltonian terms that are excluded from the CI calculation, it is also important to check that the results are reasonable. This is done by looking at the error introduced by the perturbation theory relative to the CI calculation with the same correlation model.

Figure 4.16 shows the error using first-order perturbation theory introduces to the Breit (calculated using a DC wavefunction) and vacuum polarization (calculated using a DCB wavefunction) contributions. The error is relatively small, only a few per cent for Breit, and even less for the vacuum polarization. At the multireference level, the error is exactly zero, since the multireference space in this case only contains a single CSF, and therefore the exact result is identical to the perturbative result.

However, with increasing nuclear charge, the error for Breit term becomes larger and larger. Similarly, increasing the size of the many-body space increases the error. This strongly suggests that including the Breit term in the CI matrix is highly desirable. One would also expect that in systems with more electrons and a larger CI space, the error would also be larger. The error for vacuum polarization, on the other hand, is relatively unaffected by correlation.

Figure 4.17 shows the perturbative error for the different self-energy implementations, and the order of the error is roughly the same as for Breit and vacuum polarization. Interestingly, QEDMOD seems to be slightly less amenable to perturbative treatment, with a roughly quadratic behaviour in the relative error, whereas the Flambam-Ginges and Pyykkö-Zhao potentials have a linear behaviour in Z. Finally, the hydrogenic estimates are almost unaffected by the perturbation theory, which is consistent with it effectively being a shift on the diagonal of the CI matrix.

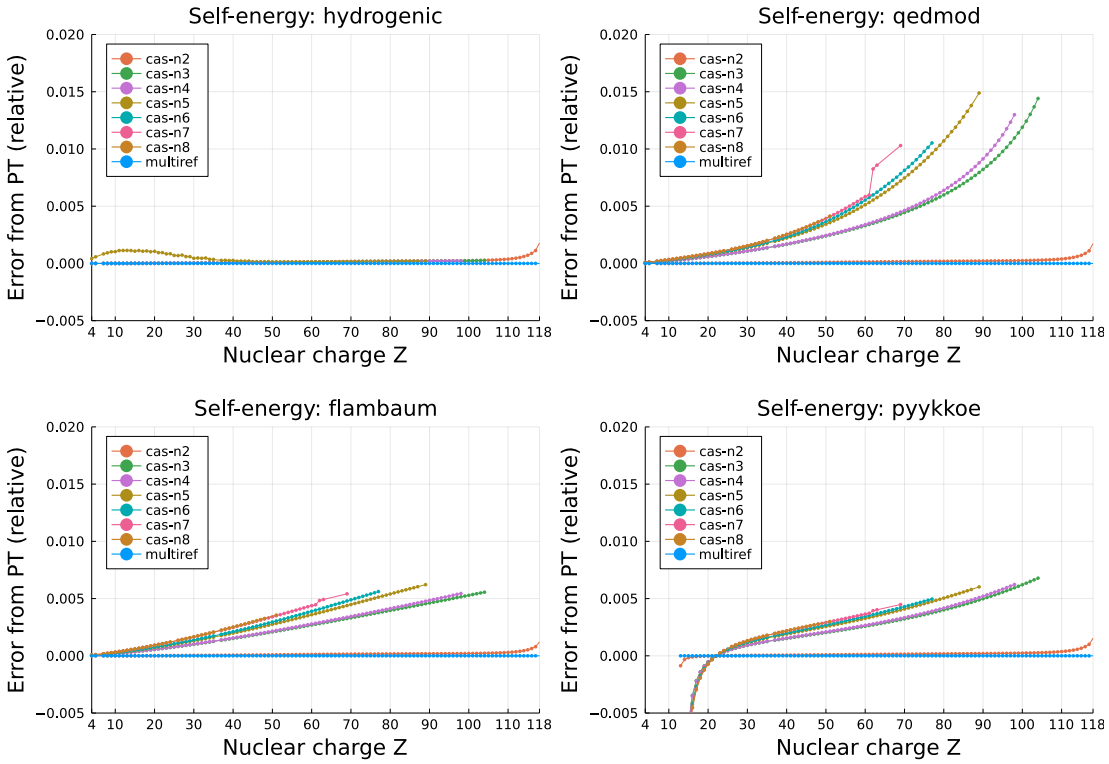

Figure 4.17: Difference in first-order perturbation theory results from the exact contributions, as a fraction of the latter, for the self-energy implementations at different correlation levels.

4.4 Implementation

The Fortran code GRASP has a long history, starting with the MCDF program package by Grant and co-workers in the middle of 1970s [153], and has gone through many iterations and versions over the decades [98, 118, 119, 120, 121, 22]. The core contribution of the work in this chapter was modifying the CI part of the code — the `rci` program — to improve the treatment of QED effects within a correlated framework. The starting point of the coding work was the MPI version of the `rci` program from the 2018 version of GRASP [22].[11]

Working on the GRASP code has presented quite a few challenges. It is an old code that has grown substantially since the 1980s, but has not been significantly refactored. While the 2018 GRASP is written in the Fortran 90 (F90) dialect[12] — all the older versions were Fortran 77

[11] Technically, as the 2018 version of GRASP was only released after the work described in this thesis started, much of the initial implementation of the QED routines was done with a modified Malmö University version of the GRASP2k package [120, 121] — the official version of GRASP preceding the 2018 version. However, transferring the routines from GRASP2k to GRASP2018 was a relatively straightforward process, as the general structure of the code had not changed.

[12] One huge downside of Fortran in the modern programming environment is the lack of resources for constructive advice on the internet relative to other programming languages. As an example, while there definitely are Fortran questions and answers on StackOverflow, there are far fewer of them compared to other more popular languages in use, such as JavaScript, Python, C or C++. As such, it is worth mentioning the book "Modern Fortran Explained" by Metcalf, Reid and Cohen [154] which offers an excellent review of modern Fortran programming practices, and discusses features of the language all the way to the Fortran 2008 dialect.

(F77), and the translation from F77 to F90 was the major contribution of the 2018 release — technically, the code is still very much in the style and spirit of Fortran 77. For a codebase the size and complexity of GRASP, this is a major drawback.

A case in point: the routines in the main GRASP codebase are not very organized. Each routine, even a tiny trivial one, lives in its own separate file. Actually, as a consequence of the automatic translation, the routines usually each come with two files — the routine, written in F77 style, and an automatically generated interface file, such that the routine could be included in other code via a `use` statement. Furthermore, the only organizational units in the code are the programs and a few "libraries" — really just collections of reused routines — and neither offer much in terms of logically organizing the code. Furthermore, the routines within each library and program simply form a flat hierarchy without any explicit structure.

Another drawback of the current codebase is the extensive use of global state, which comes from the widespread use of deprecated `common` blocks in Fortran 77 programs. The 2018 version of GRASP technically no longer uses `common` blocks but, in practice, they are still there: all the global variables of the original F77 code now live in `*_C` modules, the only purpose of which is to collect those variables. While having the variables in a module does mean that making it is no longer necessary to manually check that each routine declares the global variables in a consistent way — this is taken care of by the `.mod` file — that is actually not the main problem with the architecture. Rather, it leads to hard-to-read and hard-to-maintain code which is brittle and difficult to use, as each routine called can have unknown side effects. Unless one knows exactly how each routine interacts with every other routine via the global state, which is essentially impossible for a code the size and complexity of GRASP, and definitely not possible for someone just trying to modify one small corner of it, it becomes very hard to reason about the behaviour of the code when modifying or extending GRASP. This is also exacerbated by the lack of documentation of the routines — it is generally necessary to read all the code, and all the routines called by that code to understand what a particular part of the code is actually doing.

Both of these issues can be alleviated with judicious use of modules, a feature of Fortran added in F90. Routines should be logically organized into modules, grouped by their functionality. Global state should generally be avoided and, for any complex data, derived types should be used instead, which could be instantiated with the necessary values. The routines should then accept the data as arguments rather than storing it as global state. As a stepping stone in a gradual refactoring of the whole codebase, it may make sense to move the global state around into modules that also contain the routines that access and modify the associated global state — while not ideal, it would reduce the cognitive load of figuring out the side effects each routine causes.

While it was out of scope to rewrite the whole RCI program, where possible the existing code, and also the new code, was organized into such modules, to ease maintainability. This section, therefore, functions as a brief review and documentation for the new and modified code.

Open code and collaborative development

One major step forward for the GRASP codebase that happened, partially as part of this work, was the shift to a more open collaboration model. The authoritative version of the 2018 version of GRASP is now tracked with version control in a Git repository hosted on GitHub.[13] Rather than sending patches and tarballs around via email, changes by the collaborators are now integrated into the code on that platform via Pull Requests, which allow for much easier code review and makes the process more transparent for all collaborators. Similarly, the Issue Tracker is being

[13] As of now, the code is hosted at `https://github.com/compas/grasp`

used to keep track of bugs, ideas for improvements, and offers a single, public place for discussions on the code.

Hosting the code on a modern code collaboration platform also allows for Continuous Integration to be used, via which each and every revision of the code gets automatically tested to make sure it compiles, and also passes the limited set of automated tests. Also, as part of that process, the documentation for the codebase, generated by the Doxygen tool, gets deployed online automatically on every change, almost immediately available to collaborators and users.

Nevertheless, while the use of modern tools like that does help to maintain the code more efficiently, it ultimately boils down to the number of human hours available for it. GRASP is a complex piece of software and valuable to a large community, but there are very few resources to properly maintain it. Similarly, there are many other software packages that are adjacent to GRASP, and rather than treating them as "competitors", there should be a strong push to merge these efforts on an international level to produce a product that would be superior to any one in particular. All that requires dedicated effort, however, and there should be more focus on funding software development and code maintenance — the field of atomic physics would surely benefit from that.

Refactoring of the RCI program

The QED corrections in GRASP happen at the CI step, which is performed by the `rci` and `rci_mpi` programs. Most of the code between the two programs is actually duplicated — the primary difference between them is that one implements multiprocessing via MPI and the other is the original serial version of the code.

As the modifications to the `rci` code are extensive, they were implemented in a separate repository that just focuses on refactoring and extending the `rci` program.[14] The code was organized in a way that it is easy to compile it together with the core GRASP 2018 program, as it does use the library routines from there.

Obviously, trying to develop and maintain two almost identical codebases in parallel is a code maintenance nightmare. Because of that, the modifications were incorporated into the MPI version of the code, i.e. the 2018 GRASP version of `rci_mpi` was used as the starting point. Preferring the MPI version over the serial version was easy — the MPI implementation in the CI program is very valuable, offering a large speed boost for larger systems when running on a multicore computer. However, it can be desirable to have the serial version of the code available as well, and as part of future work, a good solution for this would be to compile both versions from the same code by guarding all the MPI-specific parts of the code with a preprocessor directive.

The first most obvious code update relative to the original version of `rci_mpi` is that the routines have been organized into groups. The `src/` directory contains the following files and directories:

rci-qed.f90 This contains the main Fortran program that runs the CI calculation, corresponding to the old executable. It gets compiled into a `rci-qed` executable.

pt.f90 This gets compiled into the `rci-qed.pt` executable and can run a follow-up calculation to the CI calculation. In addition to the standard input files, it also tries to find the TOML file, which contains all the settings the user picked in the CI calculation, in order to correctly calculate and organize the corrections.

breit Contains all the routines related to the Breit interaction.

[14] All the code, and also the development history, is available on GitHub: `https://github.com/compas/grasp-rci-qed`

libgrasp Contains various generic helper routines that are used across the program, such as the ones handling command line input, or reading-writing the various GRASP files. In the future, these will hopefully become part of the core GRASP libraries, as this would also be useful to the other GRASP programs.

mass_shifts Routines that evaluate the normal and special mass shifts.

matrixelements Refactored code for calculating the many-particle matrix elements that are shared by the main `rci-qed` program and the `rci-qed.pt` program.

qed Routines related to the QED corrections, including both the new code (e.g. new self-energy operators and related machinery) and the old routines (e.g. vacuum polarization).

rci Most of the old routines of the `rci-qed` program that are not shared with any other program.

toml An MIT-licensed C library for reading and writing TOML files [15], and the necessary Fortran wrapper module that can be used to call the library from Fortran code.

While the a more hierarchical structure was created for the files, most of the routines themselves were not modified. However, the different directories do form libraries that are relatively independent from each other, with clearly defined dependency relationships. This is further ensured by compiling them as separate static libraries in the CMake build configuration. All the libraries, except for the TOML modules, also depend on the `9290` and other libraries from the main GRASP package.

Unit and integration tests. The top-level `test/` directory contains a set of test programs and scripts that verify the behaviour of the different routines. The tests were mostly only added to the routines that were significantly modified as part of the refactoring and development. Much of the reference data was calculated with the routines themselves, so it does not strictly check correctness, but it does make sure that changes to other parts of the code do not introduce unexpected changes in the outputs. As such, the coverage is not comprehensive, but it is a major step forward from the original status quo of having no tests. The tests are also being run on GitHub on every commit using a GitHub Actions setup.

Implementation of QED contributions

At the highest level, the construction of the CI matrix is managed by the `MATRIX` subroutine, under `rci/`. This, via `GENMAT`, call `SETHAM`, which adds all the enabled contributions to matrix elements. The actual evaluation of matrix elements for each of the contributions is performed by the newly developed routines in the `grasp_rciqed_cimatrixelements` module under `matrixelements/`, and also includes summing up the one- and two-particle matrix elements together with the relevant many-particle coefficients essentially evaluating Eq. 4.4. The perturbative program `rci-qed.pt` calls the `grasp_rciqed_cimatrixelements` routines directly.

For the QED self-energy matrix elements of all the different contributions, there is only a single routine `qed_se`, which takes a matrix of one-particle elements as an argument. The QED one-particle matrix elements are generated by the routines in `grasp_rciqed_qed` under `qed/`. Some of the QED routines also require initialization of some global state, which is taken care of by the `qedse` routine that populates the single-particle matrix.

[15] Available from GitHub at: `https://github.com/skystrife/cpptoml`

```
Estimate self-energy?
> y
Choose the method for QED self-energy estimation:
    1 -- Scaled hydrogenic (default)
    2 -- QEDMOD (Shabaev et al., 2015)
    3 -- Flambaum & Ginges, 2005
    4 -- Pyykkö & Zhao, 2003
> 1
Largest n quantum number for including self-energy for orbital
n should be less or equal 8
> 4
```

Figure 4.18: The command-line user interface with questions and answers for the self-energy in the updated `rci-qed` program of GRASP.

Changes to the user interface. The user interacts with the CI program via a questions-and-answers based command-line interface, which was not considerably changed with this implementation. The only difference is that the program now asks for the user to specify the method to be used for the self-energy estimation, as illustrated in Figure 4.18. The first-order perturbation theory program `rci-qed.pt` does not ask for any interactive user input but rather requires the user to specify the name of the input file on the command line.

Hydrogenic self-energy

The implementation of the self-energy correction based on hydrogenic shifts still relies strongly on the old routines, which are now stored under `qed/legacy/`. The main entry point is the `QED_SLFEN` routine, which populates an array with the self-energy corresponding to the currently loaded orbitals. The `qedse` routine then takes that array and populates the array in accordance with Eq. (4.5).

Effective potentials

The implementations for both of the effective potentials due to Pyykkö & Zhao (`grasp_rciqed_qed_pyykkoe` module) and Flambaum & Ginges (`grasp_rciqed_qed_flambaum` module), populate an array with the potential and orbitals, and then integrate it with GRASP's `QUAD` routine, yielding the matrix element. The implementation of these routines are derived from earlier work by Christian Thierfelder and Peter Schwerdtfeger [125].

The implementation of the Flambaum & Ginges potential also includes separate numerical integrations for the inner integrals for each of the grid points. The derivative of the nuclear potential is evaluated analytically for both nuclear models currently implemented in GRASP, and therefore if any new nuclear models are implemented in the future, this part of the could would also need to be modified.

QEDMOD integration

The integration of the QEDMOD operator was technically the most complicated. It includes the code published by the authors [146, 147] and licenced under the Apache license. However, the original implementation was structured as a set of programs and was not suitable for use as

a library. Hence, the code had to undergo significant refactoring to remove the need to output temporary intermediate files etc. QEDMOD also implements its own radial grid and nuclear models, and so it was necessary to make sure that those match up with the ones implemented in GRASP, including performing an interpolation onto and from the QEDMOD grid.

Unfortunately, the QEDMOD code is also written in Fortran 77 style, which means it makes heavy use of `common` blocks. Those had to be renamed to make sure that they do not clash with anything in GRASP, and the GRASP code now includes a module `grasp_rciqed_qed_qedmod_common` which an interface for all the `common` block variables and QEDMOD routines.

All the necessary code to evaluate the matrix elements, including the initialization of the QEDMOD global state, is contained in the routines in the `grasp_rciqed_qed_qedmod` module.

Refactoring of vacuum polarization routines

While this did not change any of the numerical output of GRASP, the routines to calculate vacuum polarization contributions were also significantly refactored. The updated routines are organized very similarly to the self-energy routines, where the role of the routines that deal with specific vacuum polarization contribution types is just to provide the single-particle matrix elements of the operator. The code that incorporates it into the CI matrix is generic and shared between the different routines. The aim was to make it easy to systematically extend this part of the code with higher-order vacuum polarization contributions, such as the Wichmann-Kroll contribution. In the new structure, the different contributions can be evaluated individually, which allows the first-order perturbation theory program `rci-qed.pt` to evaluate each of the contributions separately.

Julia interface

Additional code that was developed as part of this work, in order to facilitate running calculations and working with the GRASP input and output files, was `GRASP.jl`[16] — a Julia package for working with GRASP. As it stands, it currently offers Julia functions for reading in GRASP orbital and mixing files, and also implements a basic parser for reading in the GRASP configuration state function list (CSL; `*.c`) files. It also heavily relies on the published `AtomicLevels.jl` package[17], which provides various types for representing atomic physics concepts and functions to work with them.

[16] Available as an unregistered Julia package on GitHub: `https://github.com/mortenpi/GRASP.jl`

[17] Available from the Julia package manager, with code hosted at: `https://github.com/JuliaAtoms/AtomicLevels.jl/`

Chapter 5

Epilogue

Even now, after more than a hundred years of development, atomic structure theory still offers major challenges. Both the constantly improving accuracy of experiments, requiring an increasingly thorough treatment of the theory, and the desire to have a better understanding of the various fundamental problems within the current theoretical framework spur the development of novel methods and ideas to face those challenges.

The focus of this work has been on the technical and numerical aspects of computing the electronic structure and properties of atoms and making predictions from relativistic quantum electrodynamics. Chapter 3 discusses the implementation of the one-dimensional finite element method to the relativistic Dirac equation. Taking advantage of the existing well-organized HelFEM C++ code [88], it was quite straightforward to expose it as a library in the Julia programming language and implement a simple numerical solver for the equation.

The combination of a high-level programming language (Julia [101]) and a conceptually simple basis set framework means that it is quite easy to implement various eigenvalue equations and their modifications, and then run the respective calculations and compare the results. As an example, Chapter 3 shows the ease of swapping out the nuclear potential and incorporating projection operators into the equations to eliminate the free-particle negative energy states.

While it is necessary to carefully verify convergence and make sure that everything is numerically stable, it is ultimately relatively easy to achieve results at essentially floating point precision. The size of the basis can systematically be increased, both by increasing the number of elements and by increasing the order of the primitive polynomials in each element. This leads to fully converged results, up to the point where the numerical noise from floating-point rounding errors take over.[1] The only regime where convergence was an issue was when approaching the critical charges for point nuclei (approximately $Z \approx 137$ for the full Dirac equation and $Z \approx 123$ for the projected equation).

It is worth emphasizing the numerical stability of the FEM method for the Dirac equation, especially when compared to B-splines, a widely used basis set in relativistic and non-relativistic atomic theory alike. As has been reported in the literature, it appears that it is necessary to use different orders of polynomials as the basis for the large and small components in the Dirac equation in order to obtain numerically stable results [89]. This is confirmed by anecdotal evidence from this work, as the very first attempt to write a numerical solver for the Dirac equation as part of this work was based on B-splines and led to a number of numerical difficulties. Ultimately, it was easier to switch to a different basis set rather than trying to sort out the issues themselves.

[1] To address the latter point, it would be interesting to provide an implementation of the basis which would use operators with arbitrary precision floating-point numbers, or would use interval or ball arithmetic.

While the code for the FEM method, including both the underlying HelFEM C++ library and the Julia interface, is publicly available under an open-source license, the same does not currently hold for the part of the code that actually constructs and solves the physical equations. In its current experimental proof-of-concept form, the code would be of limited use to others, and it would need to be refactored and documented before publishing. One major improvement over the current implementation would be to employ the `ContinuumArrays` framework for representing linear operators and basis sets. Using the general interface would make both HelFEM available as a basis set for other Julia projects and, similarly, it would allow other basis sets to be used directly with the existing Dirac solver code. On a related note, it would also be helpful to separate the one-dimensional FEM parts of HelFEM out of the HelFEM codebase into a separate library, that would be fully documented and designed in a way that would make it easy to provide bindings to languages like Julia or Python.

As for the physical and mathematical problems related to the Dirac equation and QED that should be explored in the future, the next major milestone would be applying the FEM method to multi-electron cases. HelFEM already implements non-relativistic Hartree-Fock and density functional theory, but preliminary attempts to use finite nuclei ran into convergence difficulties. However, this should be a technical issue and will hopefully be fixed in the future. In the future, HelFEM could also be expanded to solve the relativistic equations, and in parallel to that it would be great to include the HelFEM-powered FEM basis in the `Atoms.jl` framework[2] developed by Stefanos Carlström, which uses a slightly different formalism for performing multi-electron self-consistent field calculations in pure Julia.

Applying FEM to the Dirac equation in the high nuclear charge regime clearly shows how the bound states dive into and become degenerate and resonant with the negative energy continuum states. However, these results are purely numerical and quite qualitative, and it would be far better to find a more systematic way to isolate and quantify such bound states that are embedded in the continuum, including clearly linking them to Gamow states [15] and the work by Kuleshov and others [13, 14]. Furthermore, it would be interesting to study how those numerical resonance states could be used in multi-electron calculations.

Finally, on the topic of multi-particle calculations, it would be extremely interesting to generalize the projection operator approach. In the proper QED interpretation of the positive and negative energy states (discussed in Section 2.2), the negative energy states become associated with positive-energy positrons, and the parts of the nuclear potential operator that mix the positive and negative energy states become pair-creation and pair-annihilation terms. In the single-particle projection operator approach we are effectively removing those pair-creation and pair-annihilation terms, and so can fix the particle number. In the numerical FEM framework, it would not be too difficult to re-enable those terms and allow mixing of subspaces with different particle numbers (but known total charge), via the implementation of Fock space product states. While, in principle, this would connect us to Fock spaces with an unlimited number of particles, the energy penalty from the mass ultimately means that cutting off at relatively few particles will probably be sufficient. It would be interesting to see how this would affect the discrepancy in the predictions of the Dirac equation described in Section 3.6. An interesting special case would be solving the zero-charge case, which should lead to a form of vacuum polarization.

Chapter 4 deals with improving the treatment of QED effects in the multiconfigurational Dirac-Hartree-Fock and configuration interaction software GRASP. In order to implement the more modern approaches for multi-electron QED self-energy corrections into the CI part of GRASP, it was necessary to significantly refactor this part of the code. The implementation is tested in a multi-electron context with a four-electron Be-like system. As an incidental observa-

[2] A development version available from: `https://github.com/JuliaAtoms/Atoms.jl/`

tion from this testing was that the Be-like isoelectronic sequence exhibits a configuration change in the ground state at high nuclear charges due to the competing effects from the electron-electron interaction and finite nuclear models on the orbital energies.

In the future, and this has been discussed in the GRASP community for a long time, it would be very useful to include the additional contributions that go beyond the current Dirac-Coulomb Hamiltonian directly in the variational MCDHF part of the code. For many systems, this is necessary in order to achieve greater accuracy [20]. And while the Breit interaction would be the major contribution, QED effects such as self-energy and vacuum polarization sometimes also need to be taken into account in the variational procedure, as was demonstrated by the study of QED effects on the $^2P_{1/2}$ — $^2P_{3/2}$ transition in flourine-like systems [155, 156, 157].

More broadly, and as briefly discussed in Section 4.4, GRASP would greatly benefit from systematic maintenance work from people who have the vision and desire to implement modern software development practices. The structure of the code should be reorganized considerably, in a way that would allow for future development and maintenance to be performed more easily. It would also greatly benefit the community if the code was organized in a more accessible and hackable way, including bindings to higher-level languages such as Julia or Python, which would make exploratory coding far easier. In its current form, GRASP is quite good at doing what it does, but it is essentially impossible to improve and expand upon without an unreasonable amount of effort.

Finally, the new QED operators need to be battle-tested with interesting systems, such as various heavy and super-heavy elements. As part of this work, an attempt was made to apply the updated GRASP code on the Ni-Pd-Pt-Ds column of the periodic table, in order to predict the properties of darmstadtium (Ds). However, this project stalled because of the difficulties with the correct description of the few first low-lying states of nickel (Ni), which turned out to be a notoriously difficult problem [158, 159, 160, 161] and is still unresolved. The treatment of dense electronic spectra remains a major challenge to electronic structure theory as even large-scale multireference CI calculations and density matrix renormalization group approaches by other research groups (unpublished), performed in parallel to this work, were not able to get this spectrum right. Funnily, the spacing between the levels in palladium (Pd) and higher is actually greater, and those systems do not have this issue, so it would have been perhaps better to start from palladium instead of the lightest element in that group [162].

Bibliography

[1] Jon Grumer. "Theoretical Atomic Spectroscopy of Earthbound and Stellar Plasma". Lund University, 2016. 185 pp.

[2] P Indelicato. "QED Tests with Highly Charged Ions". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 52.23 (Dec. 14, 2019), p. 232001. ISSN: 0953-4075, 1361-6455. DOI: 10.1088/1361-6455/ab42c9.

[3] M. S. Safronova et al. "Search for New Physics with Atoms and Molecules". In: *Reviews of Modern Physics* 90.2 (June 29, 2018). ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.90.025008.

[4] Axel Beyer et al. "The Rydberg Constant and Proton Size from Atomic Hydrogen". In: *Science* 358.6359 (Oct. 6, 2017), pp. 79–85. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.aah6677.

[5] Richard P Feynman and A Zee. *QED: The Strange Theory of Light and Matter*. 2014. ISBN: 978-1-4008-4746-4.

[6] Tatsumi Aoyama et al. "Tenth-Order QED Contribution to the Electron g − 2 and an Improved Value of the Fine Structure Constant". In: *Physical Review Letters* 109.11 (Sept. 13, 2012), p. 111807. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.109.111807.

[7] D. Hanneke, S. Fogwell, and G. Gabrielse. "New Measurement of the Electron Magnetic Moment and the Fine Structure Constant". In: *Physical Review Letters* 100.12 (Mar. 26, 2008), p. 120801. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.100.120801.

[8] A Zee. *Quantum Field Theory in a Nutshell*. 2010. ISBN: 978-1-4008-3532-4.

[9] Asao Arai. *Analysis on Fock Spaces and Mathematical Theory of Quantum Fields: An Introduction to Mathematical Analysis of Quantum Fields*. New Jersey: World Scientific, 2018. 862 pp. ISBN: 978-981-320-711-0.

[10] Rudolf Haag. "On Quantum Field Theories". In: *Matematisk-fysiske Meddelser* 29.12 (1955).

[11] Bernd Thaller. *The Dirac Equation*. Texts and Monographs in Physics. Berlin ; New York: Springer-Verlag, 1992. 357 pp. ISBN: 978-3-540-54883-6 978-0-387-54883-8.

[12] Philip Ball. "Column: The Crucible". In: *Chemistry World* (Oct. 28, 2010).

[13] V M Kuleshov et al. "Coulomb Problem for a $ Z>Z_{\rm Cr}$ Nucleus". In: *Physics-Uspekhi* 58.8 (Aug. 31, 2015), pp. 785–791. ISSN: 1063-7869, 1468-4780. DOI: 10.3367/UFNe.0185.201508d.0845.

[14] S. I. Godunov, B. Machet, and M. I. Vysotsky. "Resonances in Positron Scattering on a Supercritical Nucleus and Spontaneous Production of $$e{+}e{-}$$ e + e - Pairs". In: *The European Physical Journal C* 77.11 (Nov. 2017), p. 782. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-017-5325-4.

[15] Arno Bohm, Manuel Gadella, and J. D Dollard. *Dirac Kets, Gamow Vectors and Gel'fand Triplets The Rigged Hilbert Space Formulation of Quantum Mechanics. Lectures in Mathematical Physics at the University of Texas at Austin*. Berlin: Springer Berlin, 2014. ISBN: 978-3-662-13751-2.

[16] J. Sapirstein and K. T. Cheng. "S -Matrix Calculations of Energy Levels of the Lithium Isoelectronic Sequence". In: *Physical Review A* 83.1 (Jan. 7, 2011). ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.83.012504.

[17] A. V. Malyshev et al. "QED Calculation of the Ground-State Energy of Berylliumlike Ions". In: *Physical Review A* 90.6 (Dec. 30, 2014). ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.90.062517.

[18] Pekka Pyykkö and Li-Bo Zhao. "Search for Effective Local Model Potentials for Simulation of Quantum Electrodynamic Effects in Relativistic Calculations". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 36.8 (Apr. 28, 2003), pp. 1469–1478. ISSN: 0953-4075, 1361-6455. DOI: 10.1088/0953-4075/36/8/302.

[19] V. V. Flambaum and J. S. M. Ginges. "Radiative Potential and Calculations of QED Radiative Corrections to Energy Levels and Electromagnetic Amplitudes in Many-Electron Atoms". In: *Physical Review A* 72.5 (Nov. 30, 2005), p. 052115. ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.72.052115.

[20] V. M. Shabaev, I. I. Tupitsyn, and V. A. Yerokhin. "Model Operator Approach to the Lamb Shift Calculations in Relativistic Many-Electron Atoms". In: *Physical Review A* 88.1 (July 16, 2013), p. 012513. ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.88.012513.

[21] Karol Kozioł and Gustavo A. Aucar. "QED Effects on Individual Atomic Orbital Energies". In: *The Journal of Chemical Physics* 148.13 (Apr. 7, 2018), p. 134101. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/1.5026193.

[22] C. Froese Fischer et al. "GRASP2018—A Fortran 95 Version of the General Relativistic Atomic Structure Package". In: *Computer Physics Communications* 237 (Apr. 2019), pp. 184–187. ISSN: 00104655. DOI: 10.1016/j.cpc.2018.10.032.

[23] Eite Tiesinga et al. *The 2018 CODATA Recommended Values of the Fundamental Physical Constants (Web Version 8.1)*. 2020.

[24] Michael Edward Peskin and Daniel V Schroeder. *An Introduction To Quantum Field Theory*. 1995. ISBN: 978-0-8133-4543-7 978-0-201-50397-5.

[25] Ashok Das. *Lectures on Quantum Field Theory*. Singapore ; Hackensack, NJ: World Scientific, 2008. 775 pp. ISBN: 978-981-283-286-3 978-981-283-285-6.

[26] Particle Data Group et al. "Review of Particle Physics". In: *Progress of Theoretical and Experimental Physics* 2020.8 (Aug. 14, 2020), p. 083C01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptaa104.

[27] John Ellis. "Outstanding Questions: Physics beyond the Standard Model". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 370.1961 (Feb. 28, 2012), pp. 818–830. ISSN: 1364-503X, 1471-2962. DOI: 10.1098/rsta.2011.0452.

[28] S V Troitsky. "Unsolved Problems in Particle Physics". In: (2012), p. 25.

[29] Carlo Rovelli. *Strings, Loops and Others: A Critical Survey of the Present Approaches to Quantum Gravity*. Apr. 6, 1998. arXiv: gr-qc/9803024.

[30] Carlo Rovelli. "Loop Quantum Gravity". In: (2008), p. 69.

[31] A. O. Sushkov et al. "New Experimental Limits on Non-Newtonian Forces in the Micrometer Range". In: *Physical Review Letters* 107.17 (Oct. 17, 2011), p. 171101. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.107.171101.

[32] Leslie Camilleri, Eligio Lisi, and John F. Wilkerson. "Neutrino Masses and Mixings: Status and Prospects". In: *Annual Review of Nuclear and Particle Science* 58.1 (Nov. 1, 2008), pp. 343–369. ISSN: 0163-8998, 1545-4134. DOI: 10.1146/annurev.nucl.57.090506.123038.

[33] Jonathan L. Feng. "Dark Matter Candidates from Particle Physics and Methods of Detection". In: *Annual Review of Astronomy and Astrophysics* 48.1 (Aug. 1, 2010), pp. 495–545. ISSN: 0066-4146, 1545-4282. DOI: 10.1146/annurev-astro-082708-101659.

[34] Katherine Garrett and Gintaras Duda. "Dark Matter: A Primer". In: *Advances in Astronomy* 2011 (2011), pp. 1–22. ISSN: 1687-7969, 1687-7977. DOI: 10.1155/2011/968283.

[35] Steven Weinberg. "The Cosmological Constant Problem". In: *Reviews of Modern Physics* 61.1 (Jan. 1, 1989), pp. 1–23. ISSN: 0034-6861. DOI: 10.1103/RevModPhys.61.1.

[36] David N. Spergel. "The Dark Side of Cosmology: Dark Matter and Dark Energy". In: *Science* 347.6226 (Mar. 6, 2015), pp. 1100–1102. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.aaa0980.

[37] T. Aoyama et al. "The Anomalous Magnetic Moment of the Muon in the Standard Model". In: *Physics Reports* 887 (Dec. 2020), pp. 1–166. ISSN: 03701573. DOI: 10.1016/j.physrep.2020.07.006.

[38] B. Abi et al. "Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 Ppm". In: *Physical Review Letters* 126.14 (Apr. 7, 2021), p. 141801. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.126.141801.

[39] K. Pachucki, U. D. Jentschura, and M. Pfützner. "Radiative Orbital Electron Capture by the Atomic Nucleus". In: *Physical Review C* 75.5 (May 18, 2007), p. 055502. ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.75.055502.

[40] Julian Schwinger. "Quantum Electrodynamics. I. A Covariant Formulation". In: *Physical Review* 74.10 (Nov. 15, 1948), pp. 1439–1461. ISSN: 0031-899X. DOI: 10.1103/PhysRev.74.1439.

[41] Julian Schwinger. "Quantum Electrodynamics. III. The Electromagnetic Properties of the Electron—Radiative Corrections to Scattering". In: *Physical Review* 76.6 (Sept. 15, 1949), pp. 790–817. ISSN: 0031-899X. DOI: 10.1103/PhysRev.76.790.

[42] R. P. Feynman. "Space-Time Approach to Quantum Electrodynamics". In: *Physical Review* 76.6 (Sept. 15, 1949), pp. 769–789. ISSN: 0031-899X. DOI: 10.1103/PhysRev.76.769.

[43] Willis E. Lamb and Robert C. Retherford. "Fine Structure of the Hydrogen Atom by a Microwave Method". In: *Physical Review* 72.3 (Aug. 1, 1947), pp. 241–243. ISSN: 0031-899X. DOI: 10.1103/PhysRev.72.241.

[44] H. A. Bethe. "The Electromagnetic Shift of Energy Levels". In: *Physical Review* 72.4 (Aug. 15, 1947), pp. 339–341. ISSN: 0031-899X. DOI: 10.1103/PhysRev.72.339.

[45] Norman M. Kroll and Willis E. Lamb. "On the Self-Energy of a Bound Electron". In: *Physical Review* 75.3 (Feb. 1, 1949), pp. 388–398. ISSN: 0031-899X. DOI: 10.1103/PhysRev.75.388.

[46] T.P. Gorringe and D.W. Hertzog. "Precision Muon Physics". In: *Progress in Particle and Nuclear Physics* 84 (Sept. 2015), pp. 73–123. ISSN: 01466410. DOI: 10.1016/j.ppnp.2015.06.001.

[47] R. P. Feynman. "Space-Time Approach to Non-Relativistic Quantum Mechanics". In: *Reviews of Modern Physics* 20.2 (Apr. 1, 1948), pp. 367–387. ISSN: 0034-6861. DOI: 10.1103/RevModPhys.20.367.

[48] V. Fock. "Konfigurationsraum und zweite Quantelung". In: *Zeitschrift für Physik* 75.9-10 (Sept. 1932), pp. 622–647. ISSN: 1434-6001, 1434-601X. DOI: 10.1007/BF01344458.

[49] J. M Jauch and F Rohrlich. *The Theory of Photons and Electrons: The Relativistic Quantum Field Theory of Charged Particles with Spin One-half*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1976. ISBN: 978-3-642-80951-4 978-3-642-80953-8.

[50] Milton S. Plesset. "The Dirac Electron in Simple Fields". In: *Physical Review* 41.3 (Aug. 1, 1932), pp. 278–290. ISSN: 0031-899X. DOI: 10.1103/PhysRev.41.278.

[51] Michael Ruggenthaler et al. "Quantum-Electrodynamical Density-Functional Theory: Bridging Quantum Optics and Electronic-Structure Theory". In: *Physical Review A* 90.1 (July 9, 2014), p. 012508. ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.90.012508.

[52] Michael Ruggenthaler et al. "From a Quantum-Electrodynamical Light–Matter Description to Novel Spectroscopies". In: *Nature Reviews Chemistry* 2.3 (Mar. 2018), p. 0118. ISSN: 2397-3358. DOI: 10.1038/s41570-018-0118.

[53] F. J. Dyson. "Divergence of Perturbation Theory in Quantum Electrodynamics". In: *Physical Review* 85.4 (Feb. 15, 1952), pp. 631–632. ISSN: 0031-899X. DOI: 10.1103/PhysRev.85.631.

[54] F. J. Dyson. "The Radiation Theories of Tomonaga, Schwinger, and Feynman". In: *Physical Review* 75.3 (Feb. 1, 1949), pp. 486–502. ISSN: 0031-899X. DOI: 10.1103/PhysRev.75.486.

[55] J. J. Sakurai and Jim Napolitano. *Modern Quantum Mechanics*. 2nd ed. Boston: Addison-Wesley, 2011. 550 pp. ISBN: 978-0-8053-8291-4.

[56] Murray Gell-Mann and Francis Low. "Bound States in Quantum Field Theory". In: *Physical Review* 84.2 (Oct. 15, 1951), pp. 350–354. ISSN: 0031-899X. DOI: 10.1103/PhysRev.84.350.

[57] Luca Guido Molinari. "Another Proof of Gell-Mann and Low's Theorem". In: *Journal of Mathematical Physics* 48.5 (May 2007), p. 052113. ISSN: 0022-2488, 1089-7658. DOI: 10.1063/1.2740469.

[58] J. Sucher. "S -Matrix Formalism for Level-Shift Calculations". In: *Physical Review* 107.5 (Sept. 1, 1957), pp. 1448–1449. ISSN: 0031-899X. DOI: 10.1103/PhysRev.107.1448.

[59] Peter J. Mohr, Günter Plunien, and Gerhard Soff. "QED Corrections in Heavy Atoms". In: *Physics Reports* 293.5-6 (Jan. 1998), pp. 227–369. ISSN: 03701573. DOI: 10.1016/S0370-1573(97)00046-X.

[60] G. Källén and A. Sabry. "Fourth Order Vacuum Polarization". In: *Matematisk-fysiske Meddelser* 29 (1955).

[61] W. H. Furry. "On Bound States and Scattering in Positron Theory". In: *Physical Review* 81.1 (Jan. 1, 1951), pp. 115–124. ISSN: 0031-899X. DOI: 10.1103/PhysRev.81.115.

[62] I. P Grant. *Relativistic Quantum Theory of Atoms and Molecules: Theory and Computation*. New York: Springer, 2007. ISBN: 9780387350691 9780387346717 9786610852482.

[63] G. Breit. "Dirac's Equation and the Spin-Spin Interactions of Two Electrons". In: *Physical Review* 39.4 (Feb. 15, 1932), pp. 616–624. ISSN: 0031-899X. DOI: 10.1103/PhysRev.39.616.

[64] J. A. Gaunt. "The Relativistic Theory of an Atom with Many Electrons". In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 124.793 (May 2, 1929), pp. 163–176. ISSN: 0950-1207, 2053-9150. DOI: 10.1098/rspa.1929.0104.

[65] E. A. Uehling. "Polarization Effects in the Positron Theory". In: *Physical Review* 48.1 (July 1, 1935), pp. 55–63. ISSN: 0031-899X. DOI: 10.1103/PhysRev.48.55.

[66] Eyvind H. Wichmann and Norman M. Kroll. "Vacuum Polarization in a Strong Coulomb Field". In: *Physical Review* 101.2 (Jan. 15, 1956), pp. 843–859. ISSN: 0031-899X. DOI: 10.1103/PhysRev.101.843.

[67] G. Källén. "On the Definition of the Renormalization Constants in Quantum Electrodynamics". In: *Helvetica Physica Acta* 25 (Feb. 14, 1952), p. 417.

[68] Peter J Mohr. "Self-Energy Radiative Corrections in Hydrogen-like Systems". In: *Annals of Physics* 88.1 (Nov. 1974), pp. 26–51. ISSN: 00034916. DOI: 10.1016/0003-4916(74)90398-4.

[69] Peter J Mohr. "Numerical Evaluation of the 1S12-state Radiative Level Shift". In: *Annals of Physics* 88.1 (Nov. 1974), pp. 52–87. ISSN: 00034916. DOI: 10.1016/0003-4916(74)90399-6.

[70] Peter J. Mohr. "Self-Energy of the n = 2 States in a Strong Coulomb Field". In: *Physical Review A* 26.5 (Nov. 1, 1982), pp. 2338–2354. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.26.2338.

[71] Peter J. Mohr. "Energy Levels of Hydrogen-like Atoms Predicted by Quantum Electrodynamics, 10≤Z≤40". In: *Atomic Data and Nuclear Data Tables* 29.3 (Nov. 1983), pp. 453–466. ISSN: 0092640X. DOI: 10.1016/S0092-640X(83)80002-3.

[72] Peter J. Mohr and Yong-Ki Kim. "Self-Energy of Excited States in a Strong Coulomb Field". In: *Physical Review A* 45.5 (Mar. 1, 1992), pp. 2727–2735. ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.45.2727.

[73] J.A. Lowe, C.T. Chantler, and I.P. Grant. "Self-Energy Screening Approximations in Multi-Electron Atoms". In: *Radiation Physics and Chemistry* 85 (Apr. 2013), pp. 118–123. ISSN: 0969806X. DOI: 10.1016/j.radphyschem.2013.01.004.

[74] Theodore A. Welton. "Some Observable Effects of the Quantum-Mechanical Fluctuations of the Electromagnetic Field". In: *Physical Review* 74.9 (Nov. 1, 1948), pp. 1157–1167. ISSN: 0031-899X. DOI: 10.1103/PhysRev.74.1157.

[75] Alan Robert Edmonds. *Angular Momentum in Quantum Mechanics*. Princeton Landmarks in Physics. Princeton: Princeton University Press, 1996. 145 pp. ISBN: 978-0-691-02589-6.

[76] Morris Edgar Rose. *Elementary Theory of Angular Momentum*. Dover ed. New York: Dover, 1995. 248 pp. ISBN: 978-0-486-68480-2.

[77] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskiĭ. *Quantum Theory of Angular Momentum: Irreducible Tensors, Spherical Harmonics, Vector Coupling Coefficients, 3nj Symbols*. Singapore ; Teaneck, NJ, USA: World Scientific Pub, 1988. 514 pp. ISBN: 978-9971-5-0107-5.

[78] W. Ludwig and C. Falter. *Symmetries in Physics: Group Theory Applied to Physical Problems*. 2nd extented ed. Springer Series in Solid-State Sciences 64. Berlin ; New York: Springer, 1996. 473 pp. ISBN: 978-3-540-60284-2.

[79] Xiangdong Ji, Feng Yuan, and Yong Zhao. "What We Know and What We Don't Know about the Proton Spin after 30 Years". In: *Nature Reviews Physics* 3.1 (Jan. 2021), pp. 27–38. ISSN: 2522-5820. DOI: 10.1038/s42254-020-00248-4.

[80] Peter J. Mohr, David B. Newell, and Barry N. Taylor. "CODATA Recommended Values of the Fundamental Physical Constants: 2014". In: *Reviews of Modern Physics* 88.3 (Sept. 26, 2016), p. 035009. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.88.035009.

[81] Peter Schwerdtfeger. "The Pseudopotential Approximation in Electronic Structure Theory". In: *ChemPhysChem* 12.17 (Dec. 9, 2011), pp. 3143–3155. ISSN: 14394235. DOI: 10.1002/cphc.201100387.

[82] L. F. Pašteka et al. "Atoms and Molecules in Soft Confinement Potentials". In: *Molecular Physics* 118.19-20 (Oct. 17, 2020), e1730989. ISSN: 0026-8976, 1362-3028. DOI: 10.1080/00268976.2020.1730989.

[83] Dirk Andrae. "Finite Nuclear Charge Density Distributions in Electronic Structure Calculations for Atoms and Molecules". In: *Physics Reports* 336.6 (Oct. 2000), pp. 413–525. ISSN: 03701573. DOI: 10.1016/S0370-1573(00)00007-7.

[84] I. Angeli and K.P. Marinova. "Table of Experimental Nuclear Ground State Charge Radii: An Update". In: *Atomic Data and Nuclear Data Tables* 99.1 (Jan. 2013), pp. 69–95. ISSN: 0092640X. DOI: 10.1016/j.adt.2011.12.006.

[85] W.R. Johnson and Gerhard Soff. "The Lamb Shift in Hydrogen-like Atoms, $1 \leq Z \leq 110$". In: *Atomic Data and Nuclear Data Tables* 33.3 (Nov. 1985), pp. 405–446. ISSN: 0092640X. DOI: 10.1016/0092-640X(85)90010-5.

[86] F. A. Parpia and A. K. Mohanty. "Relativistic Basis-Set Calculations for Atoms with Fermi Nuclei". In: *Physical Review A* 46.7 (Oct. 1, 1992), pp. 3735–3745. ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.46.3735.

[87] H Bachau et al. "Applications of B -Splines in Atomic and Molecular Physics". In: *Reports on Progress in Physics* 64.12 (Dec. 1, 2001), pp. 1815–1943. ISSN: 0034-4885, 1361-6633. DOI: 10.1088/0034-4885/64/12/205.

[88] Susi Lehtola. "Fully Numerical Hartree-Fock and Density Functional Calculations. I. Atoms". In: *International Journal of Quantum Chemistry* 119.19 (Oct. 5, 2019). ISSN: 0020-7608, 1097-461X. DOI: 10.1002/qua.25945.

[89] Charlotte Froese Fischer and Oleg Zatsarinny. "A B-spline Galerkin Method for the Dirac Equation". In: *Computer Physics Communications* 180.6 (June 2009), pp. 879–886. ISSN: 00104655. DOI: 10.1016/j.cpc.2008.12.010.

[90] I P Grant. "B-Spline Methods for Radial Dirac Equations". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 42.5 (Mar. 14, 2009), p. 055002. ISSN: 0953-4075, 1361-6455. DOI: 10.1088/0953-4075/42/5/055002.

[91] W. R. Johnson, S. A. Blundell, and J. Sapirstein. "Finite Basis Sets for the Dirac Equation Constructed from B Splines". In: *Physical Review A* 37.2 (Jan. 1, 1988), pp. 307–315. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.37.307.

[92] Kyle Beloy and Andrei Derevianko. "Application of the Dual-Kinetic-Balance Sets in the Relativistic Many-Body Problem of Atomic Structure". In: *Computer Physics Communications* 179.5 (Sept. 2008), pp. 310–319. ISSN: 00104655. DOI: 10.1016/j.cpc.2008.03.004.

[93] P Indelicato and JP Desclaux. *MCDFGME (ver. 2011), a multiconfiguration Dirac-Fock and general matrix elements program*. 2011.

[94] E. van Lenthe, E.J. Baerends, and J.G. Snijders. "Solving the Dirac Equation, Using the Large Component Only, in a Dirac-type Slater Orbital Basis Set". In: *Chemical Physics Letters* 236.3 (Apr. 1995), pp. 235–241. ISSN: 00092614. DOI: 10.1016/0009-2614(95)00177-6.

[95] Uzi Kaldor and Ephraim Eliav. "High-Accuracy Calculations for Heavy and Super-Heavy Elements". In: *Advances in Quantum Chemistry*. Vol. 31. Elsevier, 1998, pp. 313–336. ISBN: 978-0-12-034831-2. DOI: 10.1016/S0065-3276(08)60194-X.

[96] Osamu Matsuoka and Yoshihiro Watanabe. "An Atomic Dirac–Fock–Roothaan Program". In: *Computer Physics Communications* 139.2 (Sept. 2001), pp. 218–234. ISSN: 00104655. DOI: 10.1016/S0010-4655(01)00198-9.

[97] Sonjoy Majumder et al. "Relativistic Coupled Cluster Calculations Using Hybrid Basis Functions". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 34.23 (Dec. 14, 2001), pp. 4821–4829. ISSN: 0953-4075, 1361-6455. DOI: 10.1088/0953-4075/34/23/324.

[98] K.G. Dyall et al. "GRASP: A General-Purpose Relativistic Atomic Structure Program". In: *Computer Physics Communications* 55.3 (Oct. 1989), pp. 425–456. ISSN: 00104655. DOI: 10.1016/0010-4655(89)90136-7.

[99] T. N. Rescigno and C. W. McCurdy. "Numerical Grid Methods for Quantum-Mechanical Scattering Problems". In: *Physical Review A* 62.3 (Aug. 11, 2000), p. 032706. ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.62.032706.

[100] Jeff Bezanson et al. *Julia: A Fast Dynamic Language for Technical Computing*. Sept. 23, 2012. arXiv: 1209.5145 [cs].

[101] Jeff Bezanson et al. "Julia: A Fresh Approach to Numerical Computing". In: *SIAM Review* 59.1 (Jan. 2017), pp. 65–98. ISSN: 0036-1445, 1095-7200. DOI: 10.1137/141000671.

[102] Charlotte Froese Fischer and Oleg Zatsarinny. "A B-spline Galerkin Method for the Dirac Equation". In: *Computer Physics Communications* 180.6 (June 2009), pp. 879–886. ISSN: 00104655. DOI: 10.1016/j.cpc.2008.12.010.

[103] Oleg Zatsarinny and Charlotte Froese Fischer. "DBSR_HF: A B-spline Dirac–Hartree–Fock Program". In: *Computer Physics Communications* (2016), p. 17.

[104] P. A. M. Dirac. "A New Notation for Quantum Mechanics". In: *Mathematical Proceedings of the Cambridge Philosophical Society* 35.3 (July 1939), pp. 416–418. ISSN: 0305-0041, 1469-8064. DOI: 10.1017/S0305004100021162.

[105] Alan J. Davies. *The Finite Element Method: An Introduction with Partial Differential Equations*. 2nd ed. Oxford ; New York: Oxford University Press, 2011. 297 pp. ISBN: 978-0-19-960913-0.

[106] L. Ramdas Ram-Mohan. *Finite Element and Boundary Element Applications in Quantum Mechanics*. Oxford Texts in Applied and Engineering Mathematics 5. Oxford [England] ; New York: Oxford University Press, 2002. 605 pp. ISBN: 978-0-19-852521-9 978-0-19-852522-6.

[107] R. Layeghnejad, M. Zare, and R. Moazzemi. "Dirac Particle in a Spherical Scalar Potential Well". In: *Physical Review D* 84.12 (Dec. 20, 2011), p. 125026. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.84.125026.

[108] Per-Olov Löwdin. "On the Non-Orthogonality Problem Connected with the Use of Atomic Wave Functions in the Theory of Molecules and Crystals". In: *The Journal of Chemical Physics* 18.3 (Mar. 1950), pp. 365–375. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/1.1747632.

[109] Hans A. Bethe and Edwin E. Salpeter. *Quantum Mechanics of One- and Two-Electron Atoms*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1957. ISBN: 978-3-662-12871-8 978-3-662-12869-5. DOI: 10.1007/978-3-662-12869-5.

[110] William Desmond Evans, Peter Perry, and Heinz Siedentop. "The Spectrum of Relativistic One-Electron Atoms According to Bethe and Salpeter". In: *Communications in Mathematical Physics* 178.3 (July 1996), pp. 733–746. ISSN: 0010-3616, 1432-0916. DOI: 10.1007/BF02108822.

[111] G. Hardekopf and J. Sucher. "Relativistic Wave Equations in Momentum Space". In: *Physical Review A* 30.2 (Aug. 1, 1984), pp. 703–711. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.30.703.

[112] J. Sucher. "Foundations of the Relativistic Theory of Many-Electron Bound States". In: *International Journal of Quantum Chemistry* 25.1 (Jan. 1984), pp. 3–21. ISSN: 0020-7608, 1097-461X. DOI: 10.1002/qua.560250103.

[113] Kenneth G Dyall and Knut Fægri. *Introduction to Relativistic Quantum Chemistry*. Oxford: Oxford University Press, 2007. ISBN: 9780198032304 9780195140866 9781281158826 9786611158828 9781429486156.

[114] G. Hardekopf and J. Sucher. "Critical Coupling Constants for Relativistic Wave Equations and Vacuum Breakdown in Quantum Electrodynamics". In: *Physical Review A* 31.4 (Apr. 1, 1985), pp. 2020–2029. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.31.2020.

[115] C. Tix. "Lower Bound for the Ground State Energy of the No-Pair Hamiltonian". In: *Physics Letters B* 405.3-4 (July 1997), pp. 293–296. ISSN: 03702693. DOI: 10.1016/S0370-2693(97)00622-9.

[116] C. Tix. "Strict Positivity of a Relativistic Hamiltonian Due to Brown and Ravenhall". In: *Bulletin of the London Mathematical Society* 30.3 (May 1998), pp. 283–290. ISSN: 00246093. DOI: 10.1112/S0024609397004256.

[117] B.J. McKenzie, I.P. Grant, and P.H. Norrington. "A Program to Calculate Transverse Breit and QED Corrections to Energy Levels in a Multiconfiguration Dirac-Fock Environment". In: *Computer Physics Communications* 21.2 (Dec. 1980), pp. 233–246. ISSN: 00104655. DOI: 10.1016/0010-4655(80)90042-9.

[118] F.A. Parpia, C.Froese Fischer, and I.P. Grant. "GRASP92: A Package for Large-Scale Relativistic Atomic Structure Calculations". In: *Computer Physics Communications* 94.2-3 (Apr. 1996), pp. 249–271. ISSN: 00104655. DOI: 10.1016/0010-4655(95)00136-0.

[119] F.A. Parpia, C. Froese Fischer, and I.P. Grant. "GRASP92: A Package for Large-Scale Relativistic Atomic Structure Calculations". In: *Computer Physics Communications* 175.11-12 (Dec. 2006), pp. 745–747. ISSN: 00104655. DOI: 10.1016/j.cpc.2006.07.021.

[120] P. Jönsson et al. "The grasp2K Relativistic Atomic Structure Package". In: *Computer Physics Communications* 177.7 (Oct. 2007), pp. 597–622. ISSN: 00104655. DOI: 10.1016/j.cpc.2007.06.002.

[121] P. Jönsson et al. "New Version: Grasp2K Relativistic Atomic Structure Package". In: *Computer Physics Communications* 184.9 (Sept. 2013), pp. 2197–2203. ISSN: 00104655. DOI: 10.1016/j.cpc.2013.02.016.

[122] P Jönsson and C Froese Fischer. "SMS92: A Program for Relativistic Isotope Shift Calculations". In: *Computer Physics Communications* 100.1-2 (Feb. 1997), pp. 81–92. ISSN: 00104655. DOI: 10.1016/S0010-4655(96)00118-X.

[123] Martin Andersson and Per Jönsson. "HFSZEEMAN—A Program for Computing Weak and Intermediate Field Fine and Hyperfine Structure Zeeman Splittings from MCDHF Wave Functions". In: *Computer Physics Communications* 178.2 (Jan. 2008), pp. 156–170. ISSN: 00104655. DOI: 10.1016/j.cpc.2007.07.014.

[124] Wenxian Li et al. "Hfszeeman95—A Program for Computing Weak and Intermediate Magnetic-Field- and Hyperfine-Induced Transition Rates". In: *Computer Physics Communications* 253 (Aug. 2020), p. 107211. ISSN: 00104655. DOI: 10.1016/j.cpc.2020.107211.

[125] C. Thierfelder and P. Schwerdtfeger. "Quantum Electrodynamic Corrections for the Valence Shell in Heavy Many-Electron Atoms". In: *Physical Review A* 82.6 (Dec. 6, 2010). ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.82.062503.

[126] Charlotte Froese Fischer, Tomas Brage, and Per Jönsson. *Computational Atomic Structure: An MCHF Approach*. Bristol, UK ; Philadelphia, Penn: Institute of Physics Publ, 1997. 279 pp. ISBN: 978-0-7503-0374-3 978-0-7503-0466-5.

[127] Charlotte Froese Fischer et al. "Advanced Multiconfiguration Methods for Complex Atoms: I. Energies and Wave Functions". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 49.18 (Sept. 28, 2016), p. 182004. ISSN: 0953-4075, 1361-6455. DOI: 10.1088/0953-4075/49/18/182004.

[128] J. C. Slater. "The Theory of Complex Spectra". In: *Physical Review* 34.10 (Nov. 15, 1929), pp. 1293–1322. ISSN: 0031-899X. DOI: 10.1103/PhysRev.34.1293.

[129] Oleg Zatsarinny. "BSR: B-spline Atomic R-matrix Codes". In: *Computer Physics Communications* 174.4 (Feb. 2006), pp. 273–356. ISSN: 00104655. DOI: 10.1016/j.cpc.2005.10.006.

[130] Giulio Racah. "Theory of Complex Spectra. I". In: *Physical Review* 61.3-4 (Feb. 1, 1942), pp. 186–197. ISSN: 0031-899X. DOI: 10.1103/PhysRev.61.186.

[131] Giulio Racah. "Theory of Complex Spectra. II". In: *Physical Review* 62.9-10 (Nov. 1, 1942), pp. 438–462. ISSN: 0031-899X. DOI: 10.1103/PhysRev.62.438.

[132] Giulio Racah. "Theory of Complex Spectra. III". In: *Physical Review* 63.9-10 (May 1, 1943), pp. 367–382. ISSN: 0031-899X. DOI: 10.1103/PhysRev.63.367.

[133] Giulio Racah. "Theory of Complex Spectra. IV". In: *Physical Review* 76.9 (Nov. 1, 1949), pp. 1352–1365. ISSN: 0031-899X. DOI: 10.1103/PhysRev.76.1352.

[134] Diego R. Alcoba et al. "Configuration Interaction Wave Functions: A Seniority Number Approach". In: *The Journal of Chemical Physics* 140.23 (June 21, 2014), p. 234103. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/1.4882881.

[135] J.P. Desclaux et al. "Computational Approaches of Relativistic Models in Quantum Chemistry". In: *Handbook of Numerical Analysis*. Vol. 10. Elsevier, 2003, pp. 453–483. ISBN: 978-0-444-51248-2. DOI: 10.1016/S1570-8659(03)10006-3.

[136] Robert Duane Cowan. *The Theory of Atomic Structure and Spectra*. Los Alamos Series in Basic and Applied Sciences 3. Berkeley: University of California Press, 1981. 731 pp. ISBN: 978-0-520-03821-9.

[137] Gediminas Gaigalas, Zenonas Rudzikas, and Charlotte Froese Fischer. "An Efficient Approach for Spin - Angular Integrations in Atomic Structure Calculations". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 30.17 (Sept. 14, 1997), pp. 3747–3771. ISSN: 0953-4075, 1361-6455. DOI: 10.1088/0953-4075/30/17/006.

[138] Gediminas Gaigalas, Stephan Fritzsche, and Zenonas Rudzikas. "Reduced Coefficients of Fractional Parentage and Matrix Elements of the Tensor W(Kqkj) in Jj-Coupling". In: *Atomic Data and Nuclear Data Tables* 76.2 (Nov. 2000), pp. 235–269. ISSN: 0092640X. DOI: 10.1006/adnd.2000.0844.

[139] Gediminas Gaigalas, Stephan Fritzsche, and Ian P Grant. "Program to Calculate Pure Angular Momentum Coefficients in Jj-Coupling". In: *Computer Physics Communications* (2001), p. 16.

[140] Giovanni Li Manni, Simon D. Smart, and Ali Alavi. "Combining the Complete Active Space Self-Consistent Field Method and the Full Configuration Interaction Quantum Monte Carlo within a Super-CI Framework, with Application to Challenging Metal-Porphyrins". In: *Journal of Chemical Theory and Computation* 12.3 (Mar. 8, 2016), pp. 1245–1258. ISSN: 1549-9618, 1549-9626. DOI: 10.1021/acs.jctc.5b01190.

[141] Eva Lindroth and Ann-Marie Mårtensson-Pendrill. "Further Analysis of the Complete Breit Interaction". In: *Physical Review A* 39.8 (Apr. 1, 1989), pp. 3794–3802. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.39.3794.

[142] F. A. Parpia, M. Tong, and C. F. Fischer. "Relativistic Calculations of Nuclear Motional Effects in Many-Electron Atoms". In: *Physical Review A* 46.7 (Oct. 1, 1992), pp. 3717–3724. ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.46.3717.

[143] S. Klarsfeld and A. Maquet. "Bethe Sums for Lamb Shift Calculations in Higher Excited States". In: *Physics Letters B* 43.3 (Feb. 1973), pp. 201–203. ISSN: 03702693. DOI: 10.1016/0370-2693(73)90268-2.

[144] Vladimir B. Beresteckij et al. *Quantum Electrodynamics*. 2. ed., reprint. Course of Theoretical Physics / L. D. Landau and E. M. Lifshitz 4. Oxford: Butterworth-Heinemann, 2008. 652 pp. ISBN: 978-0-7506-3371-0.

[145] Kenneth G. Dyall. "Communication: Spectral Representation of the Lamb Shift for Atomic and Molecular Calculations". In: *The Journal of Chemical Physics* 139.2 (July 14, 2013), p. 021103. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/1.4813483.

[146] V.M. Shabaev, I.I. Tupitsyn, and V.A. Yerokhin. "QEDMOD: Fortran Program for Calculating the Model Lamb-shift Operator". In: *Computer Physics Communications* 189 (Apr. 2015), pp. 175–181. ISSN: 00104655. DOI: 10.1016/j.cpc.2014.12.002.

[147] V.M. Shabaev, I.I. Tupitsyn, and V.A. Yerokhin. "QEDMOD: Fortran Program for Calculating the Model Lamb-shift Operator". In: *Computer Physics Communications* 223 (Feb. 2018), p. 69. ISSN: 00104655. DOI: 10.1016/j.cpc.2017.10.007.

[148] L. Wayne Fullerton and G. A. Rinker. "Accurate and Efficient Methods for the Evaluation of Vacuum-Polarization Potentials of Order Z α and Z α 2". In: *Physical Review A* 13.3 (Mar. 1, 1976), pp. 1283–1287. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.13.1283.

[149] J. Blomqvist. "Vacuum Polarization in Exotic Atoms". In: *Nuclear Physics B* 48.1 (Nov. 1972), pp. 95–103. ISSN: 05503213. DOI: 10.1016/0550-3213(72)90051-X.

[150] A G Fainshtein, N L Manakov, and A A Nekipelov. "Vacuum Polarization by a Coulomb Field. Analytical Approximation of the Polarization Potential". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 24.3 (Feb. 14, 1991), pp. 559–569. ISSN: 0953-4075, 1361-6455. DOI: 10.1088/0953-4075/24/3/012.

[151] Yu. Ts. Oganessian et al. "Synthesis of the Isotopes of Elements 118 and 116 in the Cf 249 and Cm 245 + Ca 48 Fusion Reactions". In: *Physical Review C* 74.4 (Oct. 9, 2006), p. 044602. ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.74.044602.

[152] Igor A. Valuev et al. "Skyrme-Type Nuclear Interaction as a Tool for Calculating the Finite-Nuclear-Size Correction to Atomic Energy Levels and the Bound-Electron g Factor". In: *Physical Review A* 101.6 (June 1, 2020), p. 062502. ISSN: 2469-9926, 2469-9934. DOI: 10.1103/PhysRevA.101.062502.

[153] I.P. Grant. "Relativistic Atomic Structure Calculations". In: *Computer Physics Communications* 17.1-2 (Apr. 1979), pp. 149–161. ISSN: 00104655. DOI: 10.1016/0010-4655(79)90077-8.

[154] Michael Metcalf et al. *Modern Fortran Explained*. Numerical Mathematics and Scientific Computation. Oxford ; New York: Oxford University Press, 2011. 488 pp. ISBN: 978-0-19-960142-4 978-0-19-960141-7.

[155] M. C. Li et al. "Proposal of Highly Accurate Tests of Breit and QED Effects in the Ground State 2 p 5 of the F-like Isoelectronic Sequence". In: *Physical Review A* 98.2 (Aug. 7, 2018). ISSN: 2469-9926, 2469-9934. DOI: 10.1103/PhysRevA.98.020502.

[156] A. V. Volotka et al. "QED Radiative Corrections to the P 1 / 2 2 − P 3 / 2 2 Fine Structure in Fluorinelike Ions". In: *Physical Review A* 100.1 (July 23, 2019), p. 010502. ISSN: 2469-9926, 2469-9934. DOI: 10.1103/PhysRevA.100.010502.

[157] V. M. Shabaev et al. "QED Corrections to the P 1 / 2 2 − P 3 / 2 2 Fine Structure in Fluorinelike Ions: Model Lamb-shift-operator Approach". In: *Physical Review A* 101.5 (May 7, 2020), p. 052502. ISSN: 2469-9926, 2469-9934. DOI: 10.1103/PhysRevA.101.052502.

[158] Charles W. Bauschlicher, Per Siegbahn, and Lars G. M. Pettersson. "The Atomic States of Nickel". In: *Theoretica Chimica Acta* 74.6 (Dec. 1988), pp. 479–491. ISSN: 0040-5744, 1432-2234. DOI: 10.1007/BF00528018.

[159] Kerstin Andersson and Björn O. Roos. "Excitation Energies in the Nickel Atom Studied with the Complete Active Space SCF Method and Second-Order Perturbation Theory". In: *Chemical Physics Letters* 191.6 (Apr. 1992), pp. 507–514. ISSN: 00092614. DOI: 10.1016/0009-2614(92)85581-T.

[160] Ulf Litzén, James W Brault, and Anne P Thorne. "Spectrum and Term System of Neutral Nickel, Ni I". In: *Physica Scripta* 47.5 (May 1, 1993), pp. 628–673. ISSN: 0031-8949, 1402-4896. DOI: 10.1088/0031-8949/47/5/004.

[161] Nikolai B. Balabanov and Kirk A. Peterson. "Basis Set Limit Electronic Excitation Energies, Ionization Potentials, and Electron Affinities for the 3d Transition Metal Atoms: Coupled Cluster and Multireference Methods". In: *The Journal of Chemical Physics* 125.7 (Aug. 21, 2006), p. 074110. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/1.2335444.

[162] B. G. C. Lackenby, V. A. Dzuba, and V. V. Flambaum. "Calculation of Atomic Properties of Superheavy Elements Z = 110 – 112 and Their Ions". In: *Physical Review A* 101.1 (Jan. 21, 2020), p. 012514. ISSN: 2469-9926, 2469-9934. DOI: 10.1103/PhysRevA.101.012514.

[163] International Organization for Standardization. *ISO 80000–2 : Quantities and Units - Part 2: Mathematical Signs and Symbols to Be Used in the Natural Sciences and Technology*. ISO, 2009.

Appendix A

Mathematical notation

This appendix documents some of the notation and conventions used in this thesis. It primarily functions as a reference such that there would be no ambiguity in how to interpret the various mathematical expressions presented in this thesis.

A.1 Tensors and the Minkowski space

The two-dimensional Kroenecker delta δ_{ij} and the three-dimensional Levi-Civita symbol ε_{ijk}, useful when writing down various mathematical expressions in index notation, are defined as follows

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

$$\varepsilon_{ijk} = \begin{cases} 1 & \text{if } (i,j,k) \text{ is an even permutation} \\ -1 & \text{if } (i,j,k) \text{ is an odd permutation} \\ 0 & \text{if there are any repetitions in } (i,j,k) \end{cases}$$

In relativistic theory, the various quantities such as space-time points or energy-momentum values are represented as four-dimensional vectors. Such four-vectors are denoted by standard font, e.g. x, or x^μ or x_μ if the contra- and covariant components, respectively, are required to be denoted. Often, a four-vector is split up into its time-like and space-like components (in a particular frame of reference), in which case the space-like part will be a bold 3-vector

$$x^\mu = (x^0, \boldsymbol{x})$$

Often the zeroth component has a special meaning, e.g. time for a space-time point, or energy in the case of four-momentum, and is denoted with the appropriate symbol, e.g.

$$x^\mu = (t, \boldsymbol{x}), \quad p^\mu = (E, \boldsymbol{p})$$

The **Minkowski metric** is assumed to be in the so-called "particle physics" sign convention

$$g^{\mu\nu} = g_{\mu\nu} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$$

A.2 Delta function

The (one-dimensional) Dirac delta "function" is a distribution δ defined via the relationship

$$\int \mathrm{d}x\ f(x)\delta(x-y) = f(y)$$

A property of the delta function that is particularly useful when working with plane-wave functions is its relation to what is essentially the Fourier transform of the identity function

$$\int \mathrm{d}^n\boldsymbol{x}\ \mathrm{e}^{i\boldsymbol{p}\cdot\boldsymbol{x}} = (2\pi)^n\delta^{(n)}(\boldsymbol{p})$$

with the three-dimensional special case being of particular use in this work

$$\int \mathrm{d}^3\boldsymbol{x}\ \mathrm{e}^{i\boldsymbol{p}\cdot\boldsymbol{x}} = (2\pi)^3\delta^{(3)}(\boldsymbol{p})$$

Mass shell, Lorentz invariance, and normalization. Another particularly useful relationship that comes up in relativistic quantum field theory is when the delta function is being used to constrain the four-momentum of the particle to the mass shell (i.e. $p^2 = m^2$). In that case, the integral over the four-dimensional momentum space can be reduced to two integrals over the three-momentum space via

$$\int \mathrm{d}^4p\ f(p)\delta(p^2-m^2) = \int \mathrm{d}^3\boldsymbol{p}\ \frac{1}{2E_{\boldsymbol{p}}}\left[f(E_{\boldsymbol{p}},\boldsymbol{p}) + f(-E_{\boldsymbol{p}},-\boldsymbol{p})\right]$$

where $p = (E,\boldsymbol{p})$ and $f(p)$ is a four variable function $f(p) = f(E,\boldsymbol{p})$. Such an integral is Lorenz-invariant, and this relationship is the reason why the $\frac{1}{2E_{\boldsymbol{p}}}$ factor often appears in relativistic quantum field theory when it is expressed in terms of space-like quantities. To prove the relationship, one can start by

$$\begin{aligned}\int \mathrm{d}^4p\ f(p)\delta(p^2-m^2) &= \int \mathrm{d}^3\boldsymbol{p}\int_{-\infty}^{\infty}\mathrm{d}E\ f(E,\boldsymbol{p})\delta(E^2-\boldsymbol{p}^2-m^2)\\ &= \int \mathrm{d}^3\boldsymbol{p}\left[\int_{-\infty}^{0}\mathrm{d}E\ f(E,\boldsymbol{p})\delta(E^2-\boldsymbol{p}^2-m^2) + \int_{0}^{\infty}\mathrm{d}E\ f(E,\boldsymbol{p})\delta(E^2-\boldsymbol{p}^2-m^2)\right]\\ &= \int \mathrm{d}^3\boldsymbol{p}\int_{0}^{\infty}\mathrm{d}E\ \left[f(E,\boldsymbol{p}) - f(-E,\boldsymbol{p})\right]\delta(E^2-\boldsymbol{p}^2-m^2)\end{aligned}$$

where in the last step it is also necessary to take into account that $\mathrm{d}E \to -\,\mathrm{d}E$ when $E \to -E$, leading to the minus sign. Then, with the help of the variable substitution

$$y(E) = E^2-\boldsymbol{p}^2-m^2 = E^2-E_{\boldsymbol{p}}^2 \quad\Longleftrightarrow\quad E(y) = \pm\sqrt{E_{\boldsymbol{p}}^2+y},\quad \mathrm{d}E = \frac{\mathrm{d}y}{2E(y)}$$

which leads to the integration bounds changing as $y(\infty)=\infty$, $y(0) = -\boldsymbol{p}^2-m^2 = -E_{\boldsymbol{p}}^2 \le 0$, with the delta functions setting $y=0$ which leads to $E(0) = E_{\boldsymbol{p}}$, it can be shown that

$$\begin{aligned}\int \mathrm{d}^4p\ f(p)\delta(p^2-m^2) &= \int \mathrm{d}^3\boldsymbol{p}\int_{-E_{\boldsymbol{p}}}^{\infty}\ \frac{\mathrm{d}y}{2E(y)}\left[f(E(y),\boldsymbol{p}) - f(-E(y),\boldsymbol{p})\right]\delta(y)\\ &= \int \mathrm{d}^3\boldsymbol{p}\ \frac{1}{2E_{\boldsymbol{p}}}\left[f(E_{\boldsymbol{p}},\boldsymbol{p}) - f(-E_{\boldsymbol{p}},\boldsymbol{p})\right] = \int \mathrm{d}^3\boldsymbol{p}\ \frac{1}{2E_{\boldsymbol{p}}}\left[f(E_{\boldsymbol{p}},\boldsymbol{p}) + f(-E_{\boldsymbol{p}},-\boldsymbol{p})\right]\end{aligned}$$

A.3 Dirac gamma matrices

The four-by-four gamma matrices γ^μ are defined by the anti-commutation relation

$$\{\gamma^\mu, \gamma^\nu\} = 2g^{\mu\nu}$$

and therefore do not have a unique representation. Unless stated otherwise, they are assumed to be in the Dirac (or standard) representation

$$\gamma^0 = \begin{pmatrix} \mathbb{1} & 0 \\ 0 & -\mathbb{1} \end{pmatrix}, \quad \gamma^i = \begin{pmatrix} 0 & \sigma^i \\ -\sigma^i & 0 \end{pmatrix}$$

where σ^i are the two-by-two Pauli matrices and $\mathbb{1}$ the corresponding unit matrix.

$$\sigma^1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma^2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma^3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

It is also common to combine these into three-vectors $\boldsymbol{\gamma} = (\gamma^1, \gamma^2, \gamma^3)$ and $\boldsymbol{\sigma} = (\sigma^1, \sigma^2, \sigma^3)$. The gamma matrices are also linked to the $\boldsymbol{\alpha}$ and β matrices via

$$\boldsymbol{\alpha} = \gamma^0\boldsymbol{\gamma} = \begin{pmatrix} 0 & \boldsymbol{\sigma} \\ \boldsymbol{\sigma} & 0 \end{pmatrix}, \quad \beta = \gamma_0 = \begin{pmatrix} \mathbb{1} & 0 \\ 0 & -\mathbb{1} \end{pmatrix}$$

A.4 Spherical harmonics

For spherically symmetric systems, in radial coordinates, the angular parts of the orbitals can generally be solved analytically and represented in terms of spherical harmonics $Y_\ell^m(\vartheta, \phi)$. Following the conventions of the ISO 80000-2 standard "Quantities and units — Part 2: Mathematical signs and symbols to be used in the natural sciences and technology" [163], where $\phi \in [0, 2\pi]$ and $\vartheta \in [0, \pi]$ are the angular coordinates in the standard right-handed spherical coordinate system, the spherical harmonics are the solutions so the partial differential equation

$$\frac{1}{\sin\vartheta}\frac{\partial}{\partial\vartheta}\left(\sin\vartheta\frac{\partial Y_\ell^m}{\partial\vartheta}\right) + \frac{1}{\sin^2\vartheta}\frac{\partial^2 Y_\ell^m}{\partial\phi^2} + \ell(\ell+1)Y_\ell^m = 0$$

and can be written as

$$Y_\ell^m(\vartheta, \phi) = \sqrt{\frac{2\ell+1}{4\pi}\frac{(\ell-|m|)!}{(\ell+|m|)!}}P_\ell^{|m|}(\cos\vartheta)\mathrm{e}^{im\phi}$$

where $P_n^m(z)$ are the *associated Legendre functions* and solutions $y(z) = P_n^m(z)$ to the second-order differential equation

$$(1-z^2)y'' - 2zy' + \left[n(n+1) - \frac{m^2}{1-z^2}\right]y = 0$$

for $m, n \in \mathbb{N}, m \leq n$. They can be written as

$$P_n^m(z) = (-1)^m(1-z^2)^{m/2}\frac{\mathrm{d}^m}{\mathrm{d}z^m}P_n(z)$$

where in turn for $n \in \mathbb{N}$

$$P_n(z) = \frac{1}{2^n n!}\frac{\mathrm{d}^m}{\mathrm{d}z^m}\left(z^2-1\right)^n$$

are the solutions to

$$(1-z^2)\frac{\mathrm{d}^2 P_n}{\mathrm{d}z^2} - 2z\frac{\mathrm{d}P_n}{\mathrm{d}z} + n(n+1)P_n = 0$$

Spin-spherical harmonics

When working with states that are eigenstates of total angular momentum

$$\boldsymbol{J} = \boldsymbol{L} + \boldsymbol{S}$$

i.e. the sum of orbital and spin angular momenta, as is generally the case in relativistic atomic physics, the orbitals can usually be written in terms of the two-components spin-spherical harmonics $\chi_{\kappa m}(\vartheta, \phi)$. Such angular functions can be constructed from the usual spherical harmonics with the use of Clebsch-Gordan coefficients

$$\chi_{\kappa m}(\vartheta, \phi) = \sum_{\sigma} C(j, \ell, 1/2; m, m-\sigma, \sigma) Y_{\ell}^{m-\sigma}(\vartheta, \phi) \phi_{\sigma}$$

where ϕ_1 are the two spin-1/2 eigenstates of S_z

$$\phi_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \phi_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

The j and κ quantum numbers (see also Section 2.6) are related by

$$\kappa = \mp \left(j + \frac{1}{2} \right), \quad \text{for} \quad j = \ell \pm \frac{1}{2}$$

Appendix B

Solutions of the Dirac equation

The Dirac equation for a single spin-1/2 particle in a central radial potential $V(r)$ (Eq. (2.27) in Section 2.6) can be written as

$$\left[i\gamma^\mu\partial_\mu - m + e\gamma^0 V(r)\right]\psi(x) = 0 \tag{B.1}$$

where γ^μ are the Dirac gamma matrices and ψ are 4-component complex functions on $\mathbb{R}^3$.

B.1 Free-particle solutions

The free-particle special case of the equation (B.1) without any external fields (see also Eq. 2.14 in Section 2.4) is

$$\left(i\gamma^\mu\partial_\mu - m\right)\psi(x) = 0$$

Following the derivation in Grant's book [62], this equation has plane-wave solutions of the form

$$\psi(x) = u(p)\mathrm{e}^{-ix\cdot p}$$

where $x^\mu = (t, \boldsymbol{x})$ and $p^\mu = (E, \boldsymbol{p})$ are four-vectors. The time-dependency can be further split off into an exponential

$$\psi(t, \boldsymbol{x}) = u(E, \boldsymbol{p})\mathrm{e}^{ix\cdot\boldsymbol{p}}\mathrm{e}^{-iEp}$$

and if the Dirac equation is applied to this expression, it leads to a four-by-four matrix equation for $u(p)$

$$\left(\gamma_\mu p^\mu - m\right)u(p) = 0 \tag{B.2}$$

which therefore should have four independent solutions.

Multiplying Eq. (B.2) with $(\gamma_\mu p^\mu + m)$ on the left leads to the constraint $p^\mu p_\mu = E^2 - \boldsymbol{p}^2 = m^2$ on the valid four-momenta that the solution can have, restricting them onto the so-called "mass shell". The three-momentum $\boldsymbol{p}$ can be arbitrary, and leads to energies $p^0 = E = \pm E(\boldsymbol{p})$ for each three-momentum, where $E(\boldsymbol{p}) = \sqrt{m^2 - \boldsymbol{p}^2}$ is the positive energy square root of the relativistic energy. This shows how the solutions of the free-particle Dirac equation split into two distinct subsets: the positive energy continuum with $E \geq mc^2$ and the negative energy continuum with $E \leq -mc^2$.

$u(p)$ is a four-component value which can be split up into the upper and lower component pairs

$$u(p) = \begin{pmatrix} u_1(p) \\ u_2(p) \end{pmatrix}$$

and equation (B.2) implies a relationship between $u_1(p)$ and $u_2(p)$

$$u_2(p) = \frac{\boldsymbol{\sigma}\cdot\boldsymbol{p}}{p^0+m}u_1(p) \quad \Leftrightarrow \quad u_1(p) = \frac{\boldsymbol{\sigma}\cdot\boldsymbol{p}}{p^0-m}u_2(p) \tag{B.3}$$

However, that is the only constraint and so it is possible to pick e.g. u_1 completely freely, and that then fixes u_2 (or vice verse). As u_1 (or u_2) is just a two-component complex vector, there are two linearly independent choices, e.g.

$$\phi_1 = \begin{pmatrix}1\\0\end{pmatrix}, \quad \phi_2 = \begin{pmatrix}0\\1\end{pmatrix}$$

To avoid the $p^0 = m$ singularity in the denominators in (B.3), the convention is to fix the upper component to ϕ^r if $p^0 > 0$ and the lower component when $p^0 < 0$. Altogether, for each $\boldsymbol{p}$, there are four linearly independent $u(p)$ values, and to write them down more compactly it is helpful to defined the symbol

$$\Gamma(\boldsymbol{p}) = \frac{\boldsymbol{\sigma}\cdot\boldsymbol{p}}{E(\boldsymbol{p})+m}$$

For $p^0 = E(\boldsymbol{p}) > 0$ there therefore are the two solutions

$$u(p) = u^{(r)}(\boldsymbol{p}) \propto \begin{pmatrix}\phi_r\\ \Gamma(\boldsymbol{p})\phi_r\end{pmatrix}, \quad r = 1,2$$

The other two solution when $p^0 = -E(\boldsymbol{p}) < 0$, in which case

$$u_1(p) = \frac{\boldsymbol{\sigma}\cdot\boldsymbol{p}}{p^0-m}u_2(p) = \frac{\boldsymbol{\sigma}\cdot\boldsymbol{p}}{-E(\boldsymbol{p})-m}u_2(p) = -\frac{\boldsymbol{\sigma}\cdot\boldsymbol{p}}{E(\boldsymbol{p})+m}u_2(p) = -\Gamma(\boldsymbol{p})u_2(p)$$

can be written as

$$u(p) = v^{(r)}(\boldsymbol{p}) \propto \begin{pmatrix}-\Gamma(\boldsymbol{p})\phi_r\\ \phi_r\end{pmatrix}, \quad r = 1,2$$

Normalization. The normalization for $u^{(r)}$ and $v^{(r)}$ is in terms of the inner product $\bar\psi\psi = \psi^\dagger\gamma^0\psi$. It can be shown that

$$\bar u^{(r)}(p)v^{(s)}(p) = \bar v^{(r)}(p)u^{(s)}(p) = 0, \quad \bar u^{(r)}(p)u^{(s)}(p) \propto \delta_{rs}, \quad \bar v^{(r)}(p)v^{(s)}(p) \propto \delta_{rs}$$

The Lorentz invariant normalization condition is

$$\bar u^{(r)}(p)u^{(s)}(p) = \delta_{rs} = \bar v^{(r)}(p)v^{(s)}(p)$$

which leads to the following definition of $u^{(r)}(\boldsymbol{p})$ and $v^{(r)}(\boldsymbol{p})$

$$\begin{aligned} u^{(r)}(\boldsymbol{p}) &= \sqrt{\frac{E(\boldsymbol{p})+m}{2m}}\begin{pmatrix}\phi^r\\ \Gamma(\boldsymbol{p})\phi^r\end{pmatrix}, \quad r = 1,2\\ v^{(r)}(\boldsymbol{p}) &= \sqrt{\frac{E(\boldsymbol{p})+m}{2m}}\begin{pmatrix}-\Gamma(\boldsymbol{p})\phi^r\\ \phi^r\end{pmatrix}, \quad r = 1,2 \end{aligned} \tag{B.4}$$

Finally, it can also be shown that u and v are also orthogonal in terms of the standard inner product

$$u^{(r)\dagger}(\boldsymbol{p})u^{(s)}(\boldsymbol{p}) = \frac{E_p}{m}\delta_{rs}, \quad v^{(r)\dagger}(\boldsymbol{p})v^{(s)}(\boldsymbol{p}) = -\frac{E_p}{m}\delta_{rs}$$
$$u^{(r)\dagger}(\boldsymbol{p})v^{(s)}(\boldsymbol{p}) = 0 = v^{(r)\dagger}(\boldsymbol{p})u^{(s)}(\boldsymbol{p})$$

albeit not normalized.

B.2 Bound state solutions for point charge

A special case of the external-field Dirac equation (B.1) that can be solved analytically is if the central potential is generated by a point nucleus, leading to a Coulomb potential centered on the origin

$$V(r) = -Z/r \tag{B.5}$$

The solutions are generally expressed in terms of spherical coordinates (see also Section 2.6) where the orbitals are represented in terms of radial complex single variable functions $P(r)$ and $Q(r)$ defined by Eq. (2.28)

$$\psi(\boldsymbol{x}) = \frac{1}{r}\begin{pmatrix} P(r)\chi_{\kappa m}(\theta,\varphi) \\ iQ(r)\chi_{-\kappa m}(\theta,\varphi) \end{pmatrix}$$

where $\chi_{\kappa m}(\theta,\varphi)$ are the spin-spherical harmonics (see Appendix A), leading to the radial Dirac equation

$$\begin{pmatrix} mc^2 - \frac{Z}{r} & \hbar c\left(-\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right) \\ \hbar c\left(\frac{\mathrm{d}}{\mathrm{d}r} + \frac{\kappa}{r}\right) & -mc^2 - \frac{Z}{r} \end{pmatrix} \begin{pmatrix} P(r) \\ Q(r) \end{pmatrix} = E \begin{pmatrix} P(r) \\ Q(r) \end{pmatrix} \tag{B.6}$$

which is a two-component matrix differential equation, and a special case of Equation (2.29) with the potential (B.5).

The bound state solutions for a given angular momentum subspace labelled by the κ quantum number to Eq. (B.6) can be written as [62, Eq. (3.3.22) and (3.3.23)]

$$\begin{aligned} P(r) &= \mathcal{N}\sqrt{c + E/c}\ \rho^{\gamma}\mathrm{e}^{-\rho/2}[-n_r M(-n_r+1, 2\gamma+1; \rho) + (N-\kappa)M(-n_r, 2\gamma+1; \rho)] \\ Q(r) &= \mathcal{N}\sqrt{c - E/c}\ \rho^{\gamma}\mathrm{e}^{-\rho/2}[-n_r M(-n_r+1, 2\gamma+1; \rho) - (N-\kappa)M(-n_r, 2\gamma+1; \rho)] \end{aligned}$$

where $\mathcal{N}$ is the normalization constant

$$\mathcal{N} = \sqrt{\frac{Z\alpha}{2N^2(N-\kappa)}\frac{\Gamma(2\gamma+n_r+1)}{n_r[\Gamma(2\gamma+1)]^2}}$$

$n_r = n + |\kappa|$ is the *inner quantum number*, $\gamma = \sqrt{\kappa^2 - (Z\alpha)^2} > 0$, $\rho = 2\lambda r$ and $\lambda = Z/N$ where N is the *apparent principal quantum number*.

$$N = \sqrt{n^2 - 2n_r(|\kappa| - \gamma)} = \sqrt{(n_r+\gamma)^2 + (Z\alpha)^2}$$

The functions $M(a, b; \rho)$ are the solutions to the Kummer's confluent hypergeometric equation [62, Eq. (A.3.14) and (A.3.15)].

Appendix C

Additional FEM figures

This appendix contains additional figures for Chapter 3 with additional parameters showing that the behaviours described are consistent across a range of parameter values. Specifically, there are three groups of figures:

1. Additional examples of the non-relativistic grid checks with a range of nuclear charges Z and nuclear models, complementing Figure 3.8. The figures show how well the non-relativistic $1s$ energy converges as a function of the z parameter of the exponential grid, defined by Eq. (3.4). The parameter determines how close to the origin the element boundaries are distributed, with higher z placing the elements closer to the origin. (Figures C.1, C.2, C.3, C.4)

2. Additional examples of the relativistic grid checks, complementing Figures 3.16 and 3.17. As in the non-relativistic case, they show the convergence of the relativistic $1s$ energy as a function of the grid parameter z for various nuclear charges Z and nuclear models. Figures C.5, C.7, C.8, C.9 show the convergence for different nuclear models with 6th order polynomials, whereas Figure C.6 shows the point nucleus case with 10th order polynomials.

3. Additional examples of the consistency of relativistic $1s$ energy estimates with various primitive basis polynomials orders and element distributions, complementing Figures 3.19, 3.20, and 3.21 from Section 3.4 with additional nuclear charge Z values and nuclear models. The values reported on the bottom of each plot is the finite nuclear correction values, i.e. the $1s$ energy difference from the analytical point nucleus energy. The energies and their uncertainties were determined using the algorithm described in Section 3.3. In some cases, especially when the grid points are distributed and at higher Z values, it is not always possible to achieve good convergence with different z values. The estimate 284 ± 1.6 Ha on the bottom-left plots in Figure C.10 shows that near the critical charge $Z \approx 137$ it very difficult to achieve convergence in the case of a point nucleus. (Figures C.10, C.11, C.12, C.13)

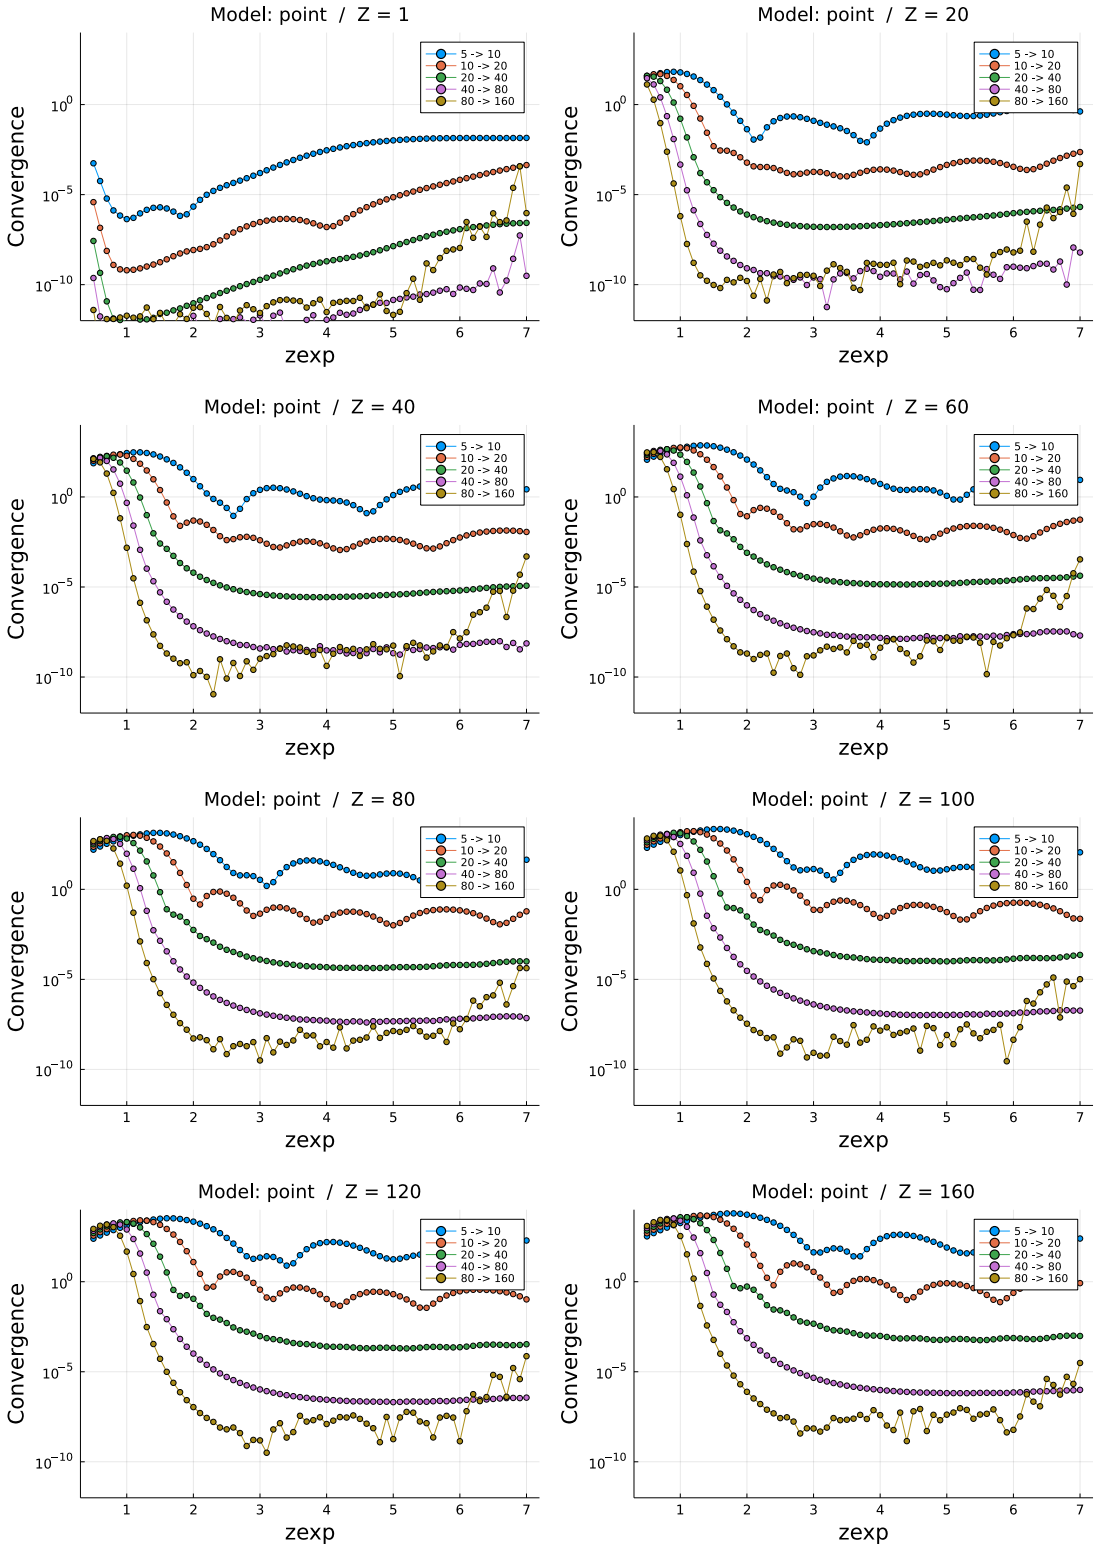

Figure C.1: Convergence of the non-relativistic $1s$ energy as a function of the z parameter of the exponential grid for the **point nucleus** for various nuclear charge Z values ranging from 1 to 160. The calculations were performed with order 6 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\rm rms} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

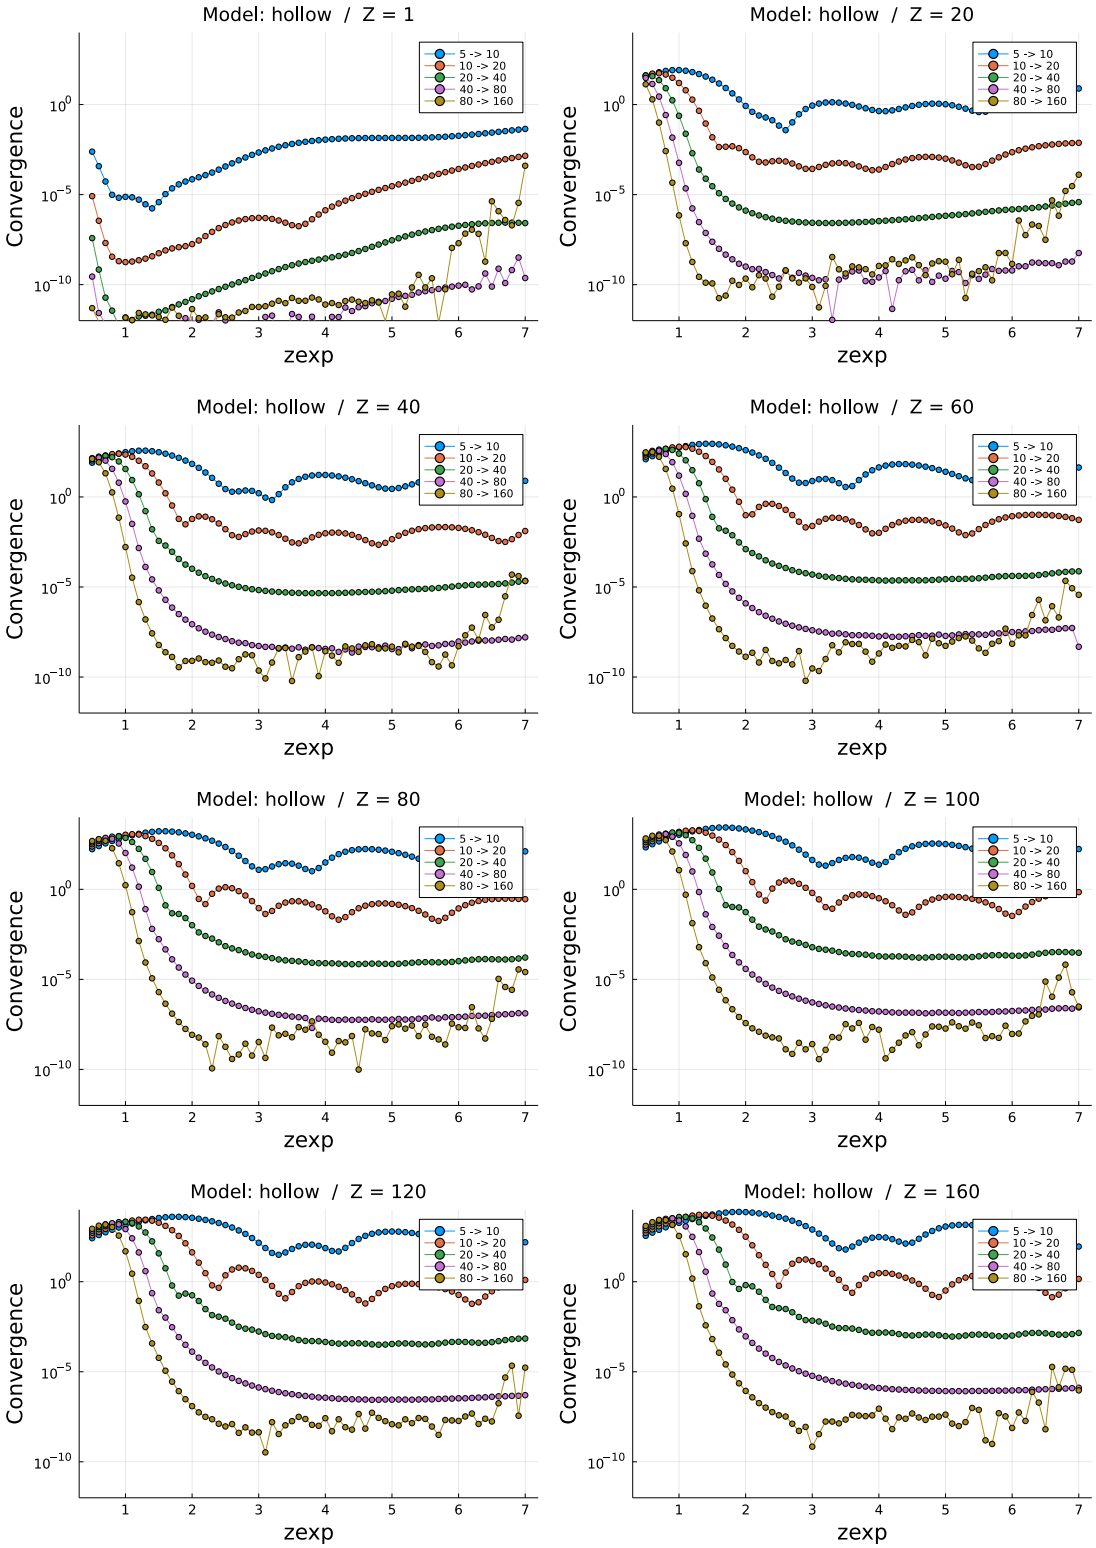

Figure C.2: Convergence of the non-relativistic $1s$ energy as a function of the z parameter of the exponential grid for the **hollow nucleus** for various nuclear charge Z values ranging from 1 to 160. The calculations were performed with order 6 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\mathrm{rms}} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

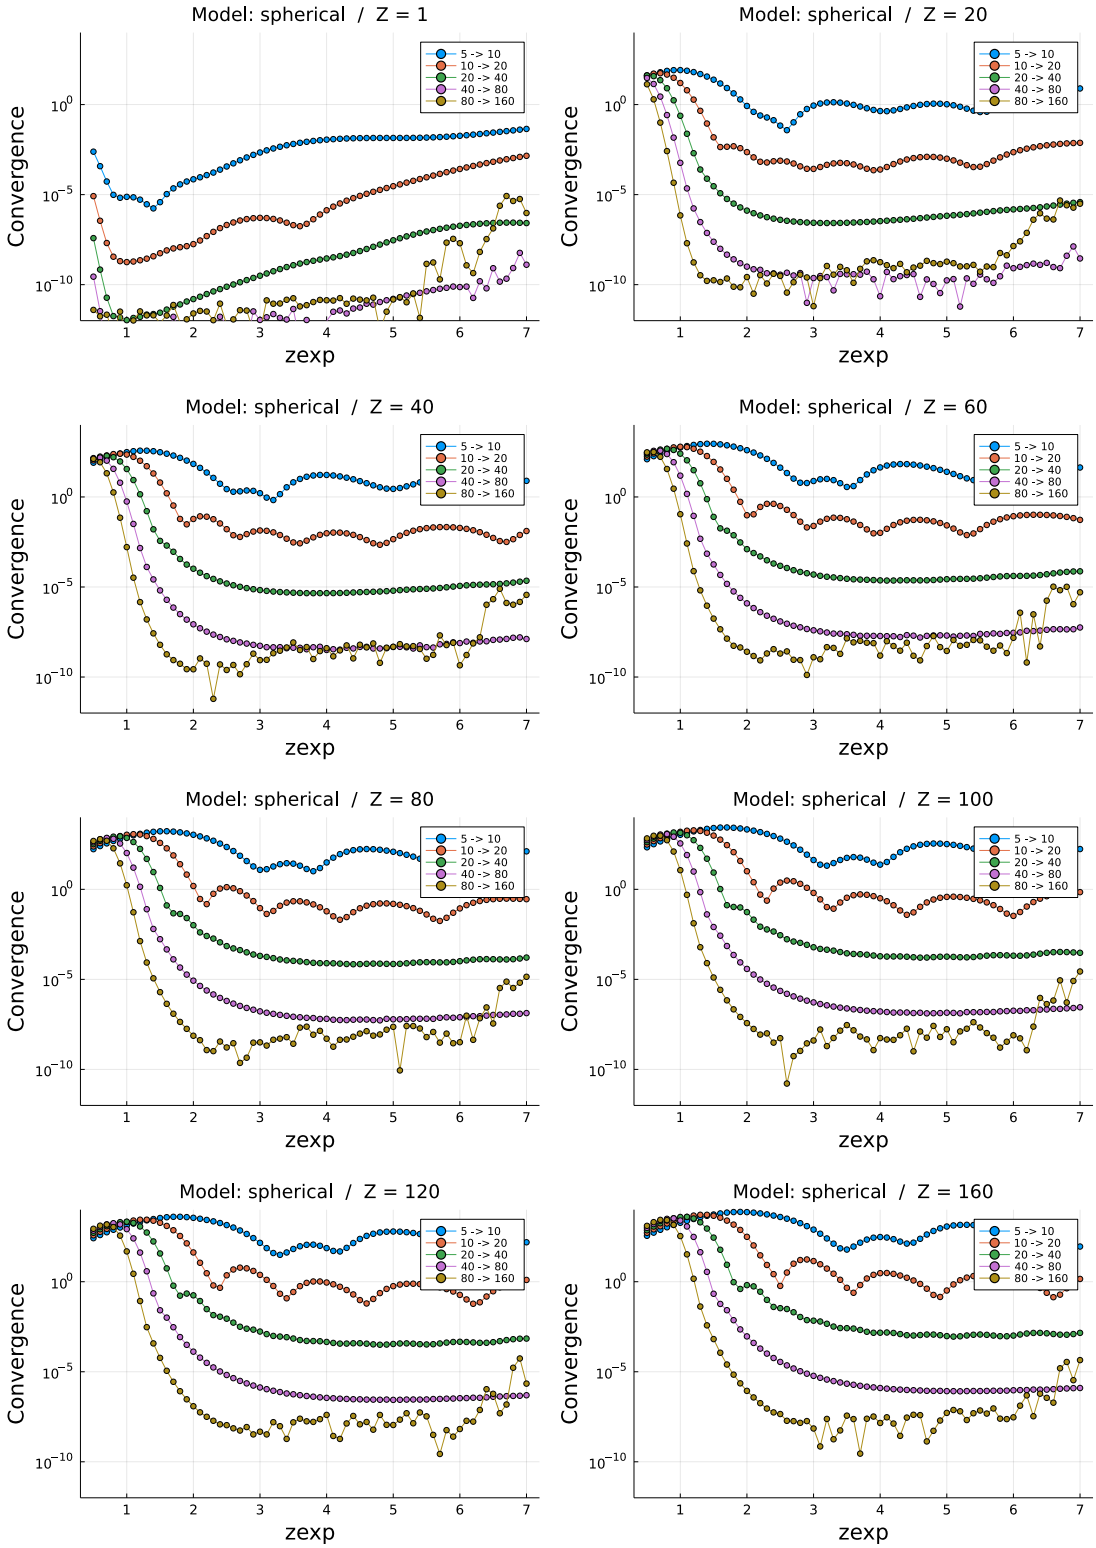

Figure C.3: Convergence of the non-relativistic $1s$ energy as a function of the z parameter of the exponential grid for the **spherical nucleus** for various nuclear charge Z values ranging from 1 to 160. The calculations were performed with order 6 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\rm rms} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

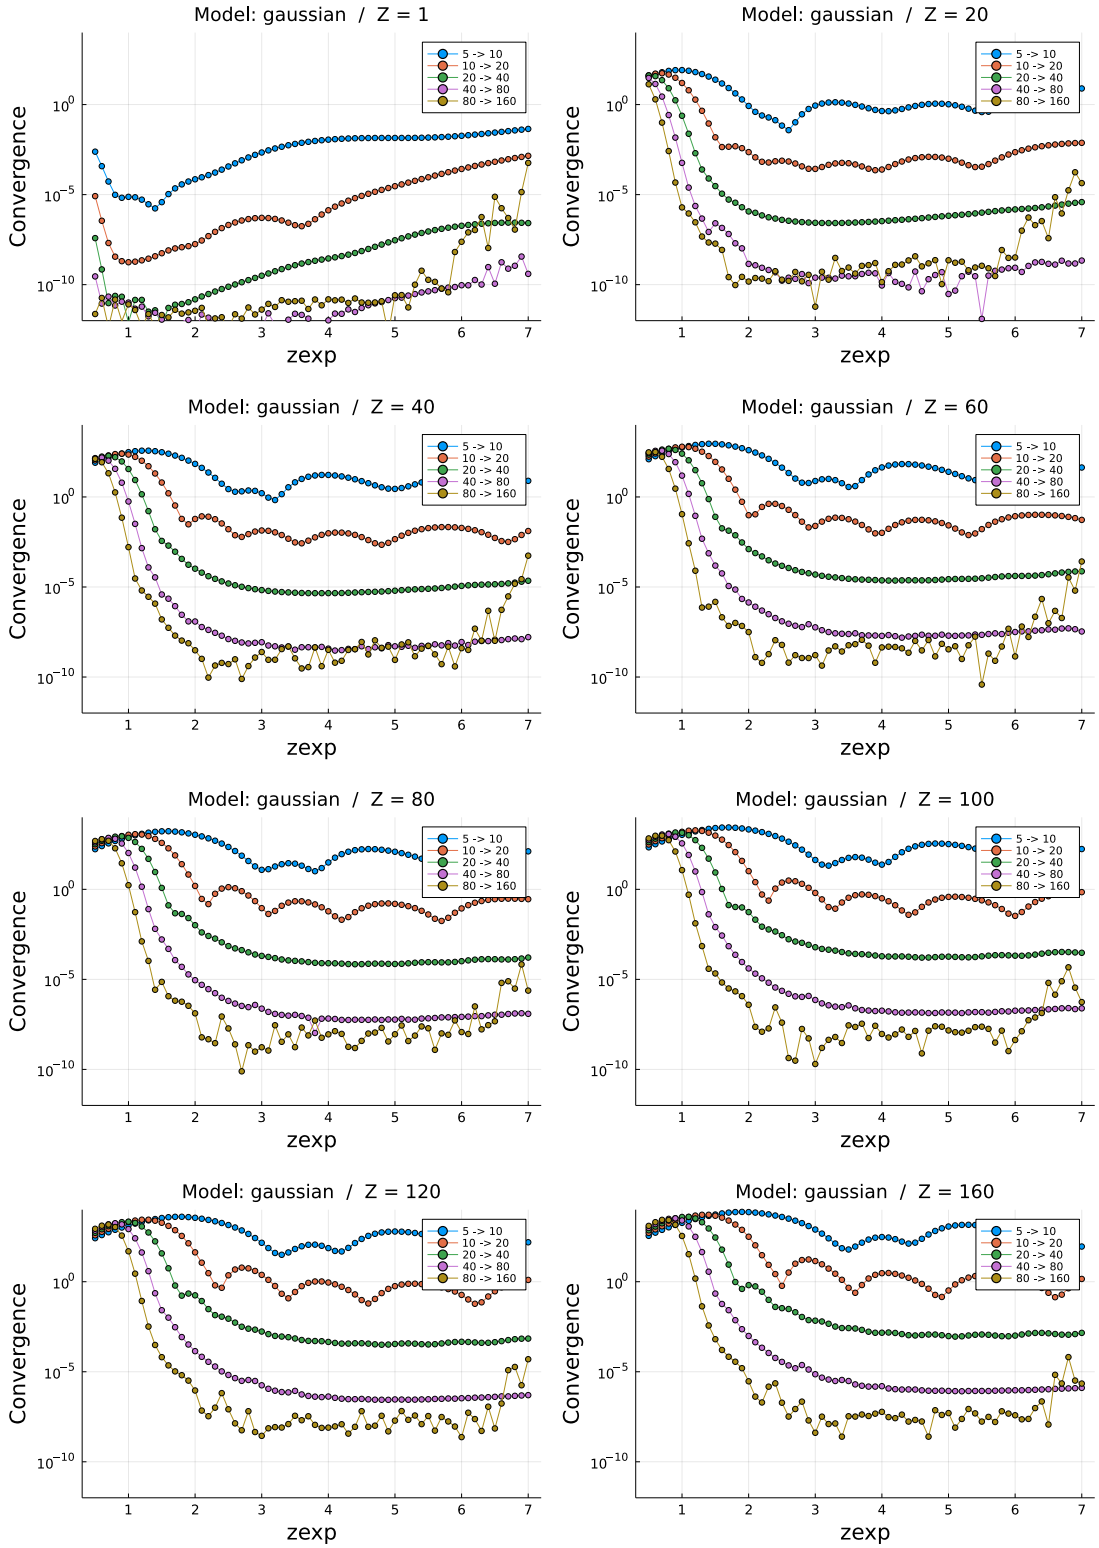

Figure C.4: Convergence of the non-relativistic $1s$ energy as a function of the z parameter of the exponential grid for the **Gaussian nucleus** for various nuclear charge Z values ranging from 1 to 160. The calculations were performed with order 6 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\mathrm{rms}} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

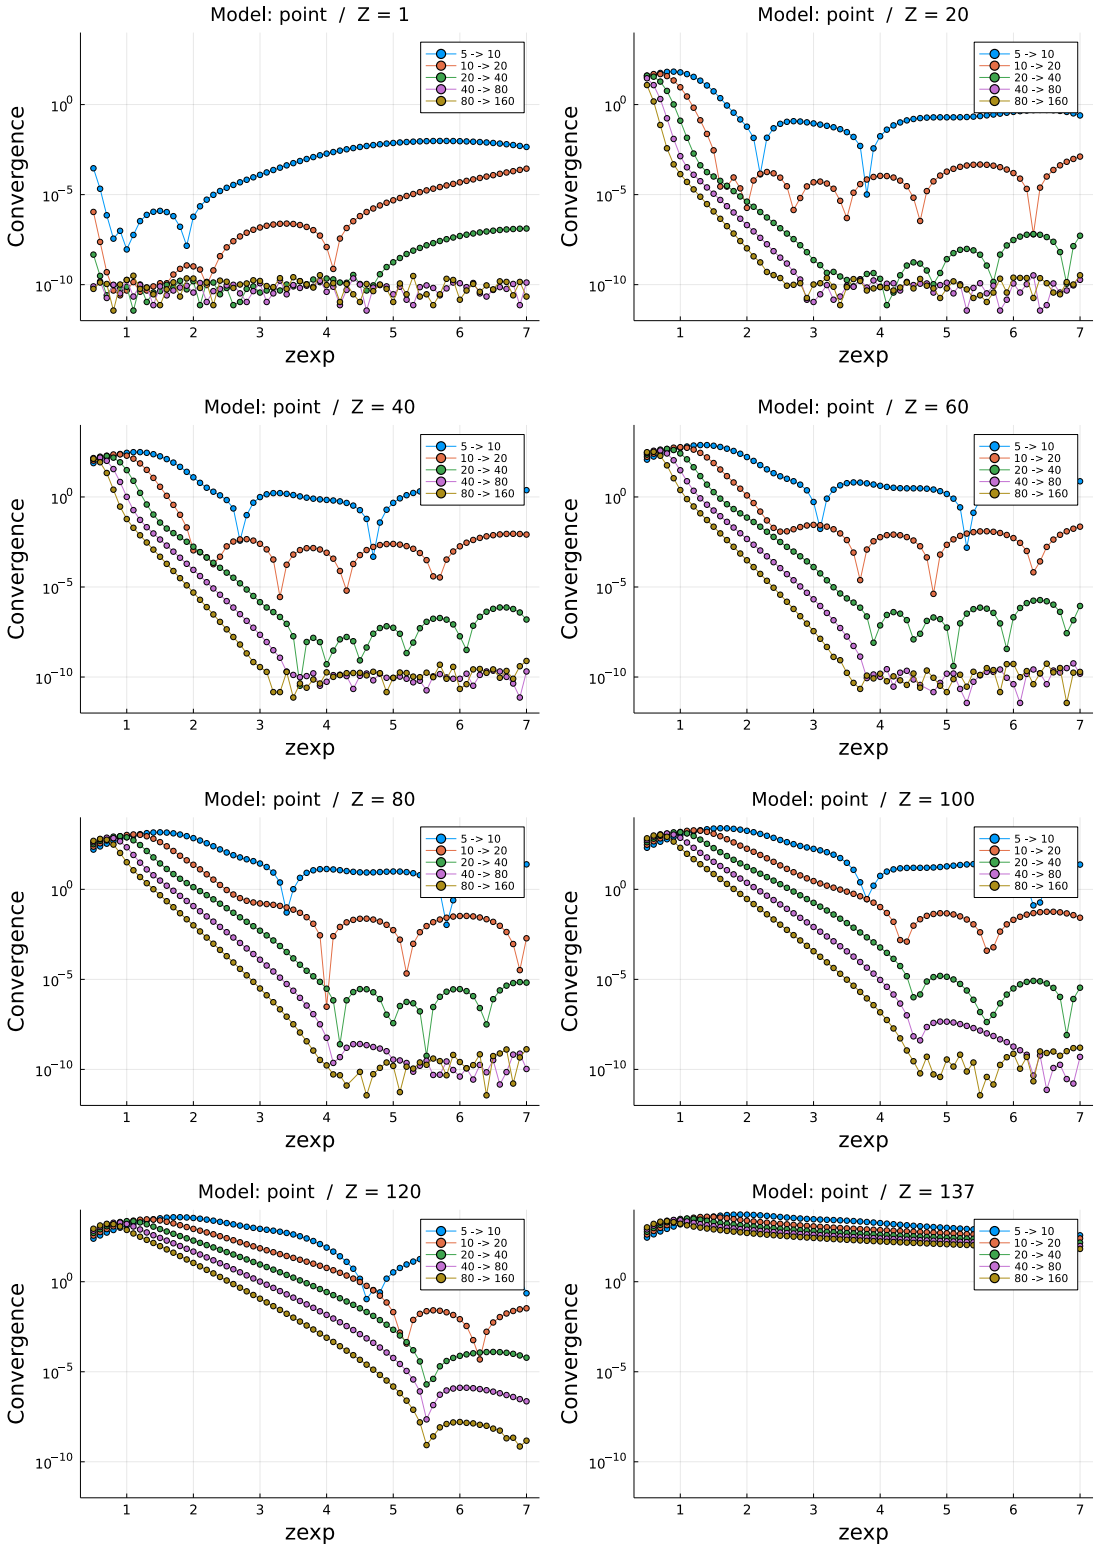

Figure C.5: Convergence of the relativistic $1s$ energy as a function of the z parameter of the exponential grid for the **point nucleus** for various nuclear charge Z values ranging from 1 to 160. The calculations were performed with order 6 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\mathrm{rms}} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

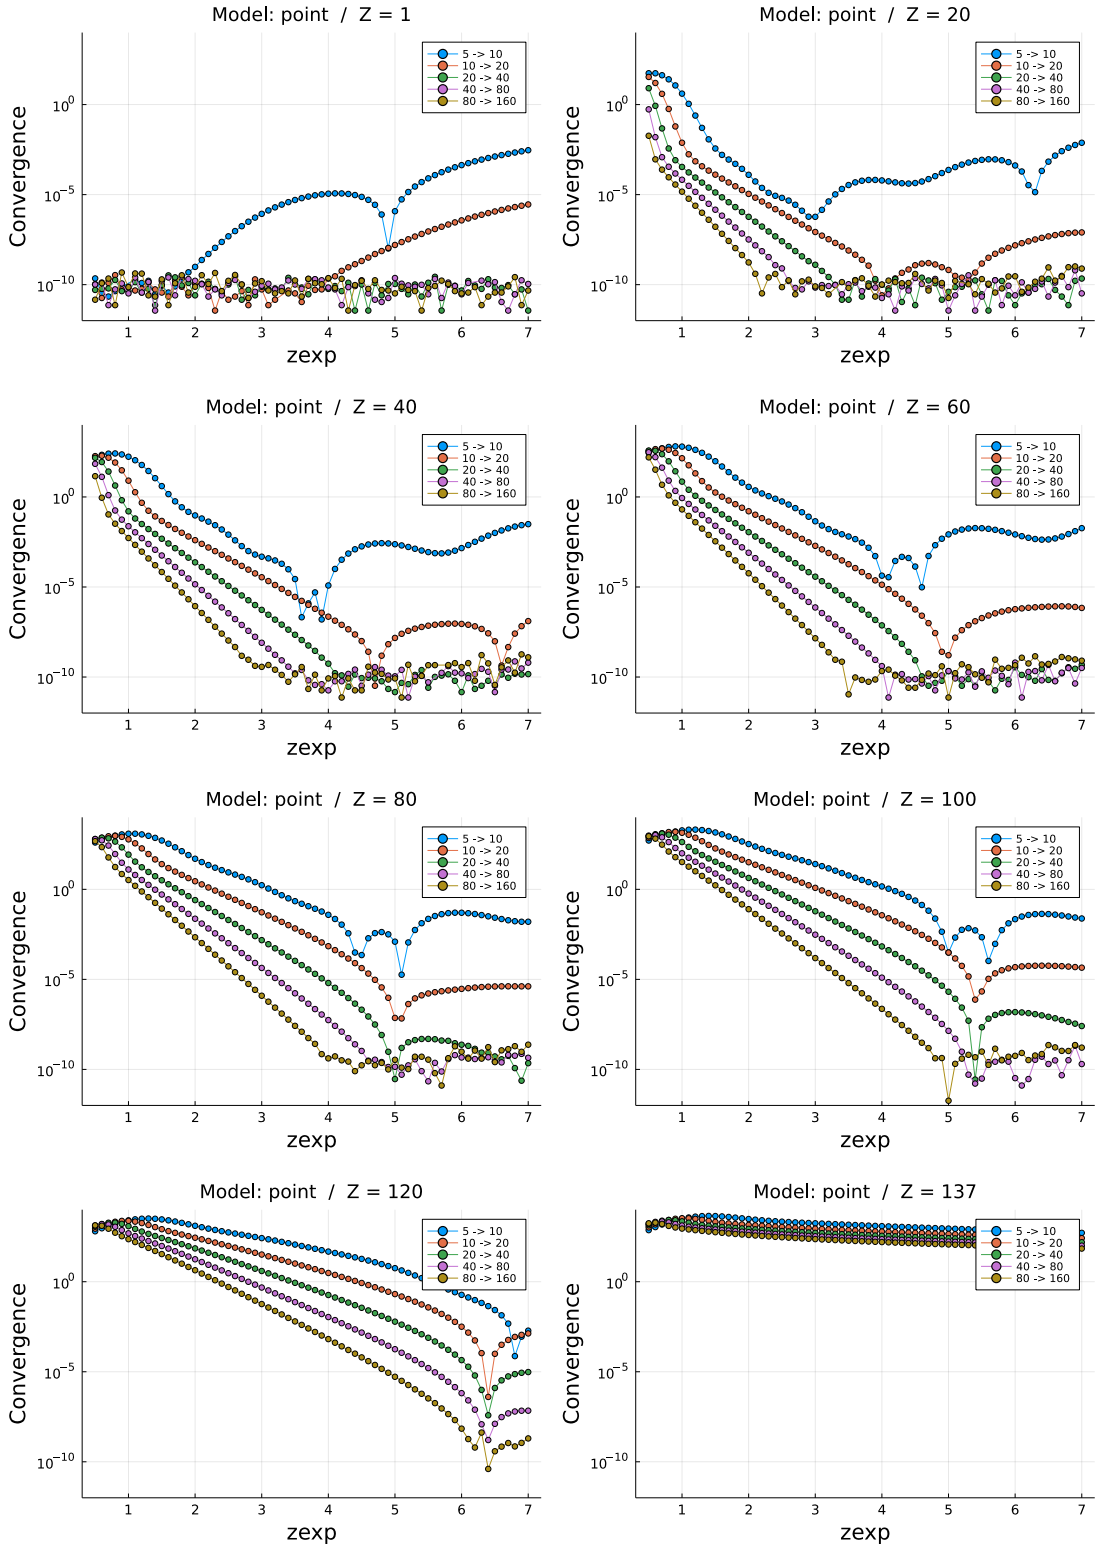

Figure C.6: Convergence of the relativistic $1s$ energy as a function of the z parameter of the exponential grid for the **point nucleus** for various nuclear charge Z values ranging from 1 to 160. The calculations were performed with order 10 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\mathrm{rms}} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

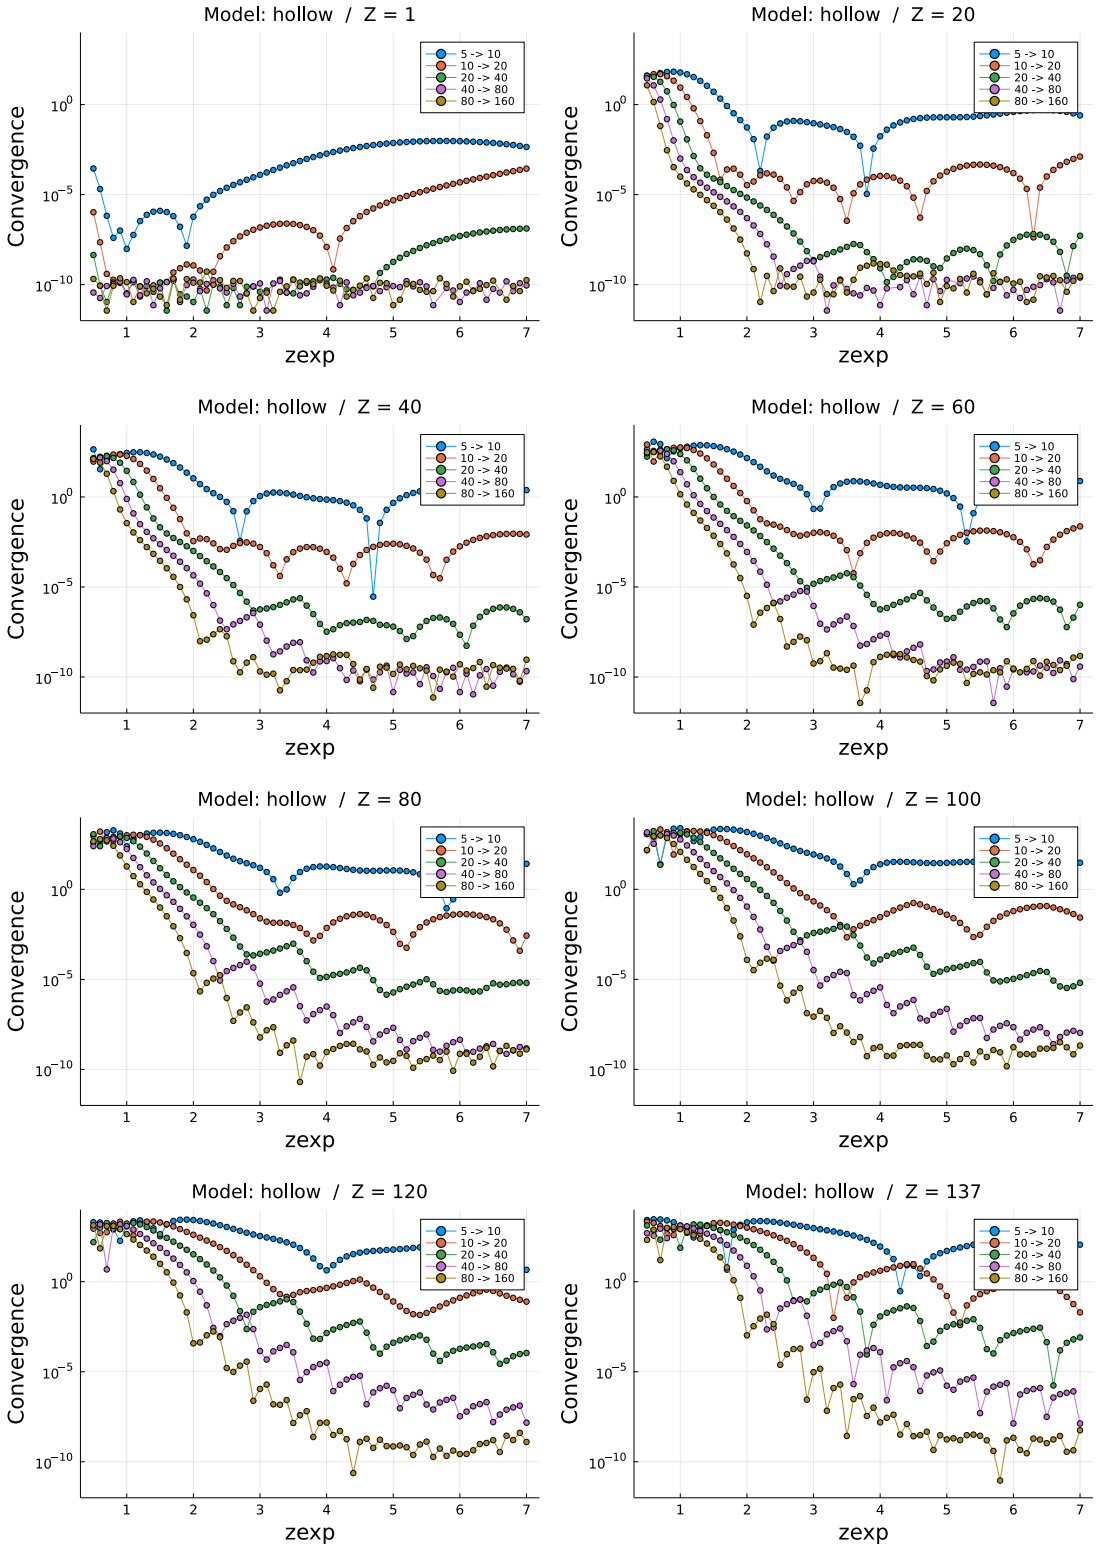

Figure C.7: Convergence of the relativistic $1s$ energy as a function of the z parameter of the exponential grid for the **hollow nucleus** for various nuclear charge Z values ranging from 1 to 160. The calculations were performed with order 6 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\mathrm{rms}} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

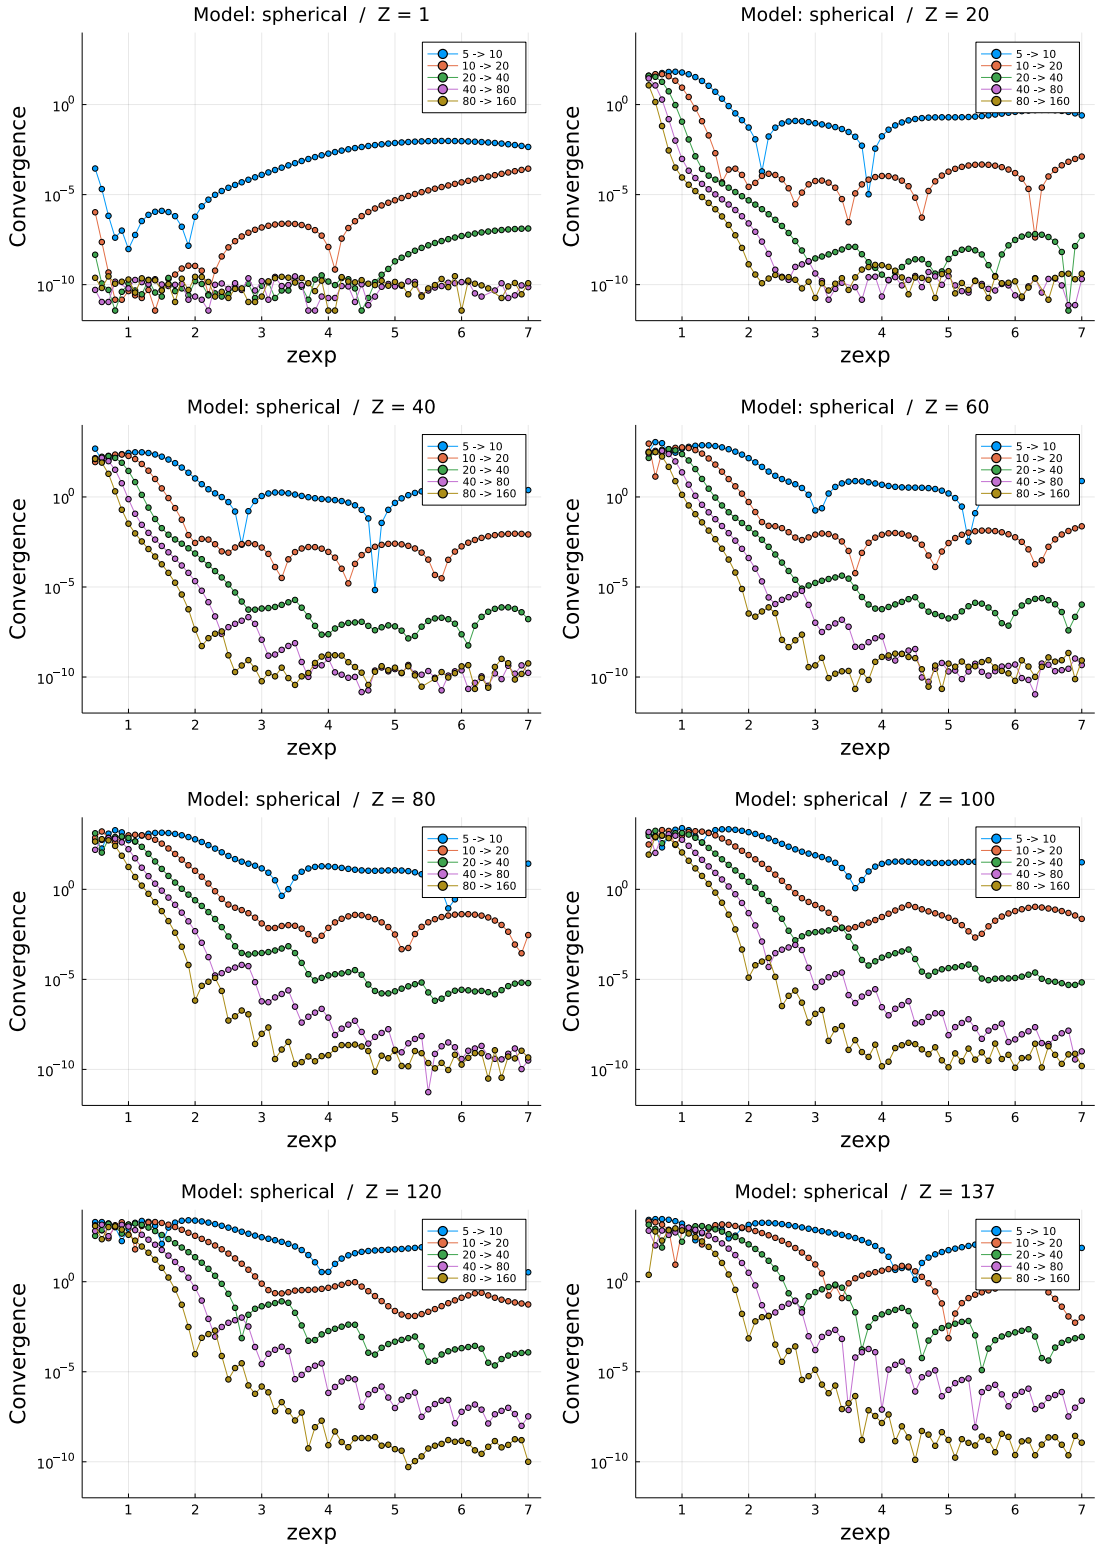

Figure C.8: Convergence of the relativistic $1s$ energy as a function of the z parameter of the exponential grid for the **spherical nucleus** for various nuclear charge Z values ranging from 1 to 160. The calculations were performed with order 6 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\mathrm{rms}} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

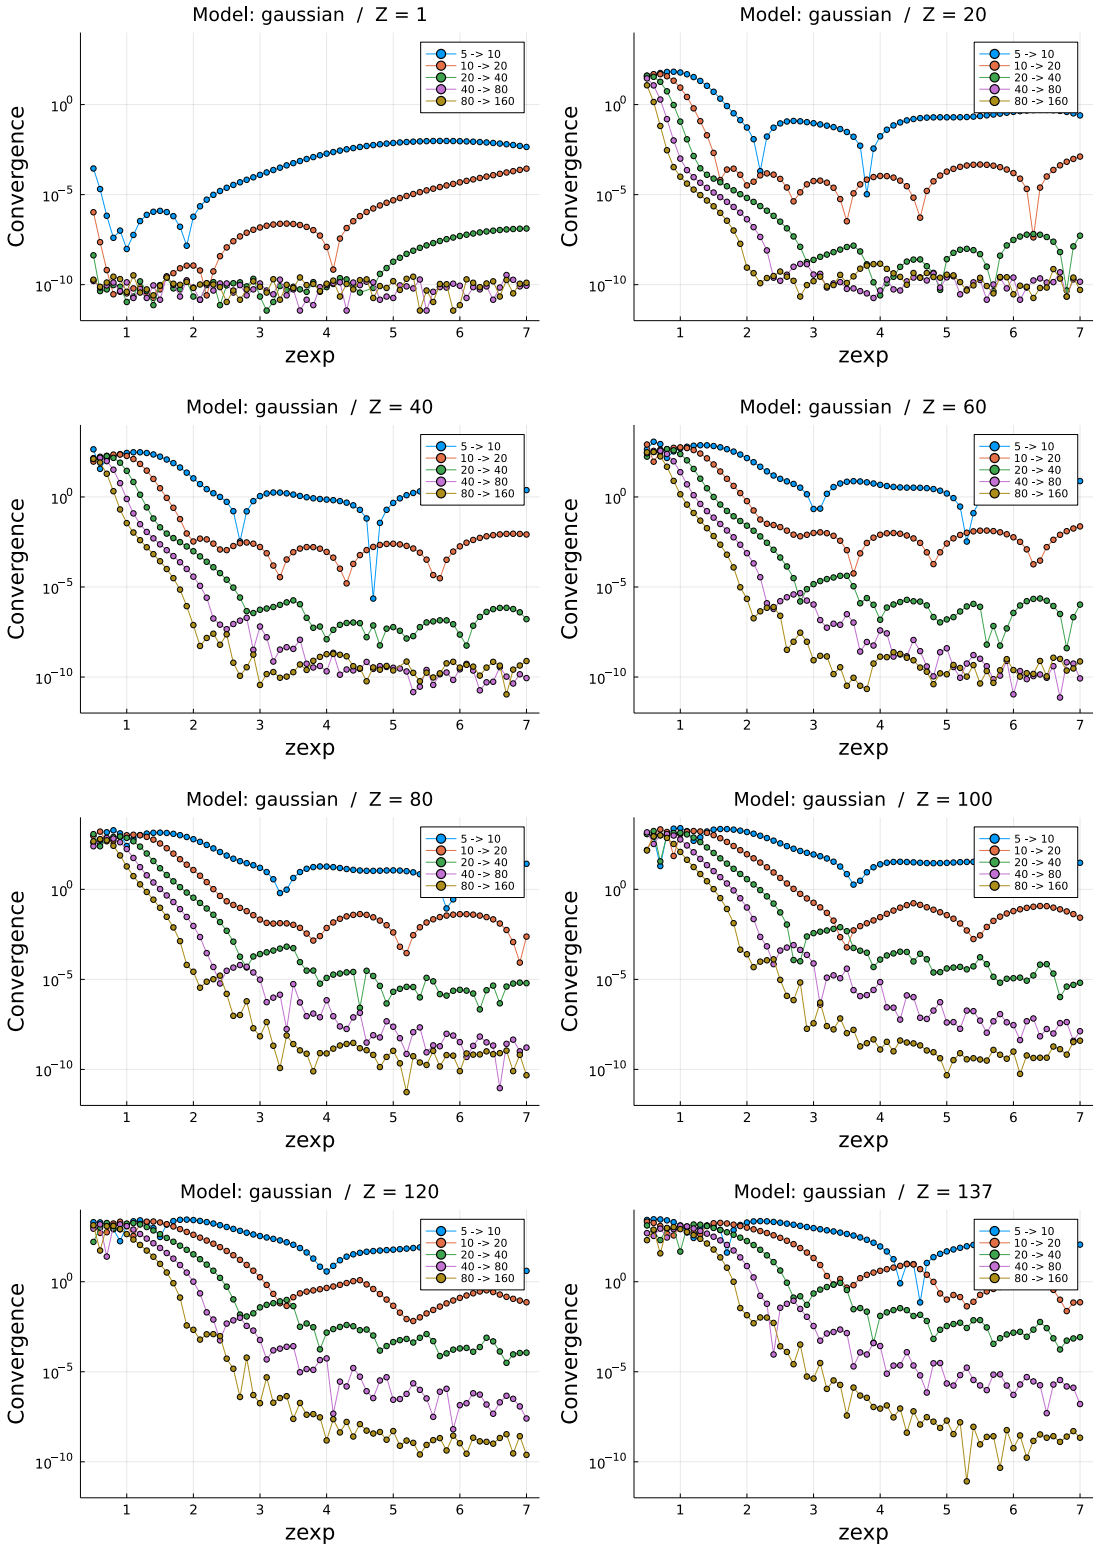

Figure C.9: Convergence of the relativistic $1s$ energy as a function of the z parameter of the exponential grid for the **Gaussian nucleus** for various nuclear charge Z values ranging from 1 to 160. The calculations were performed with order 6 polynomials, $r_{\max} = 40\ a_0$ and for nuclear models with $R_{\rm rms} = 10^{-4}\ a_0$. The convergence measure is defined as the (absolute) change in energy as the number of elements is doubled (as indicated in the legend).

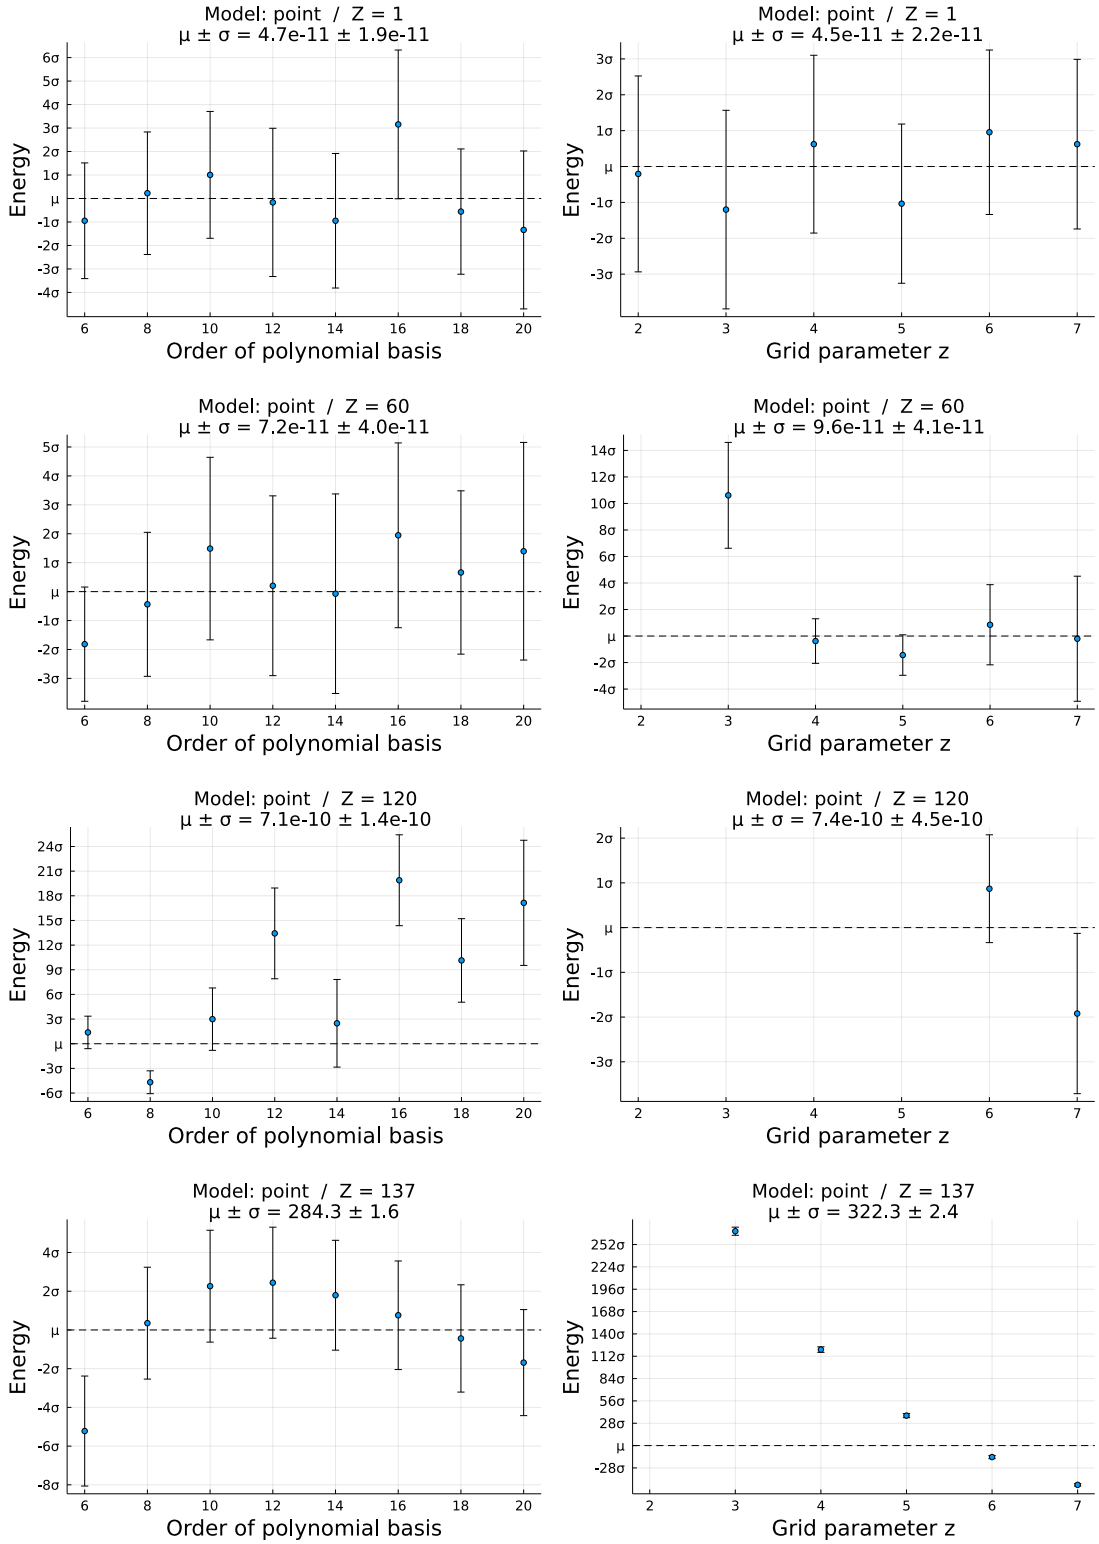

Figure C.10: Consistency of the estimates of the relativistic $1s$ energy values, together with their uncertainties for the **point nucleus**. The left column is consistency between runs with different basis polynomial orders and the right column is with different z parameter values of the exponential grid. The values on the y-axis are relative to the mean μ and uncertainty σ of the FNC correction (reported below each graph in Hartrees) of all the data points.

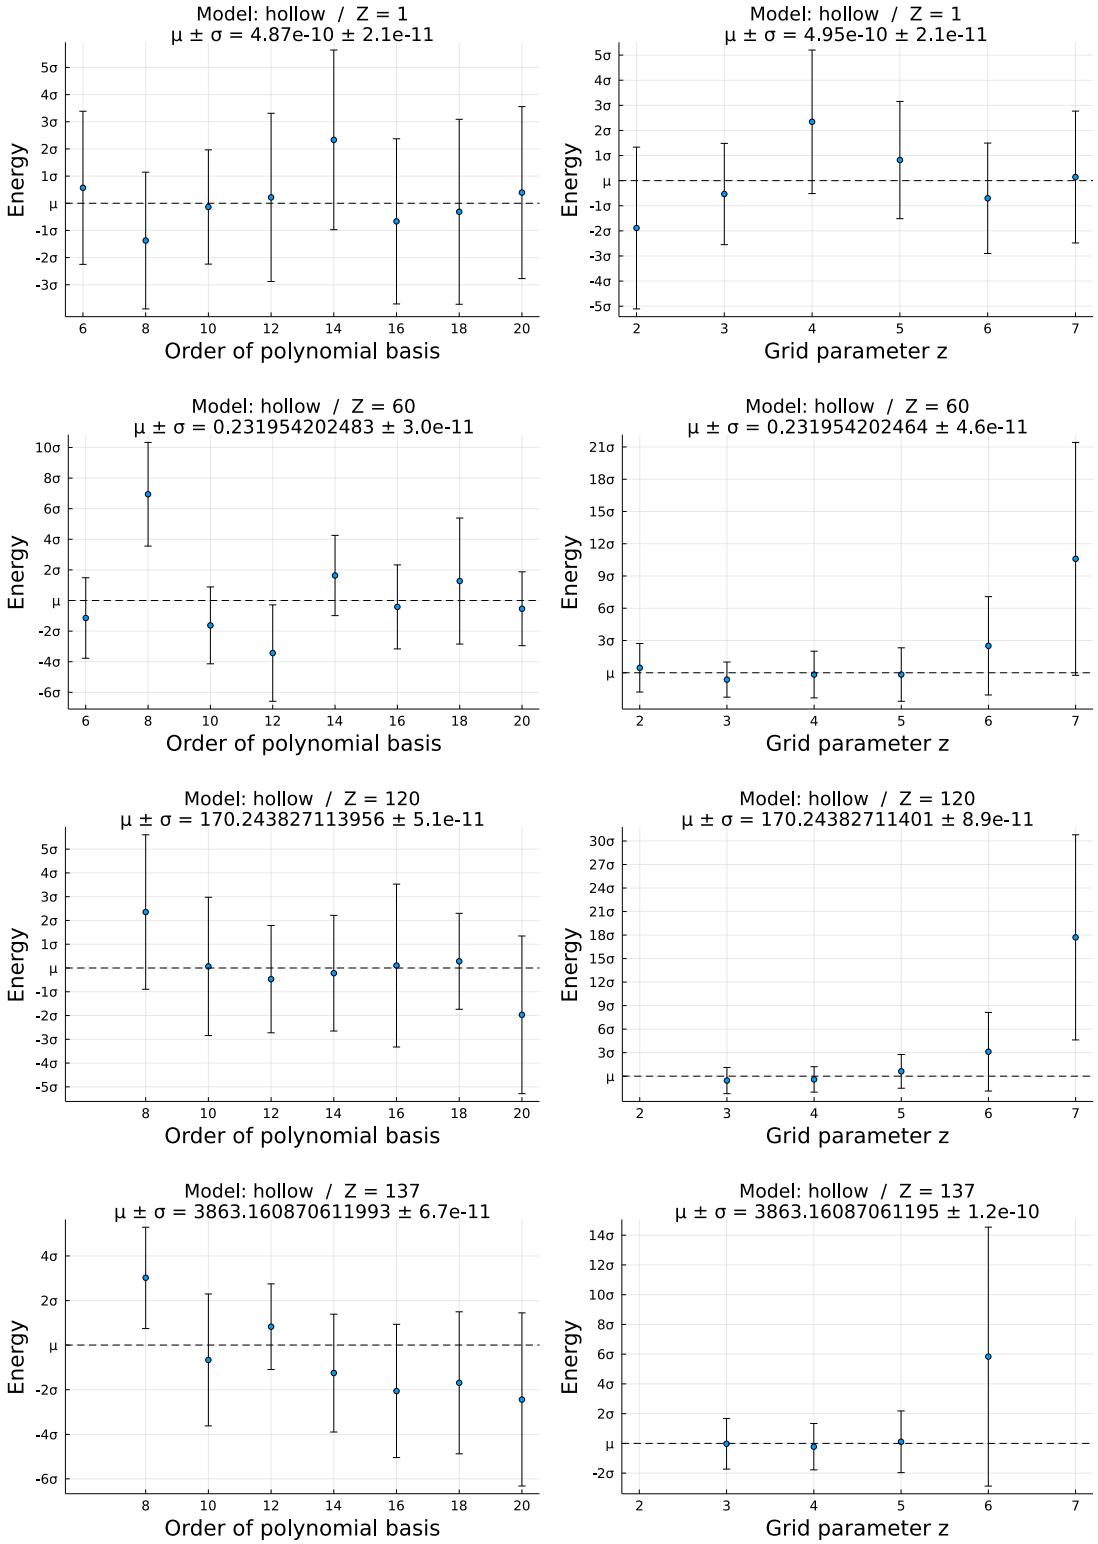

Figure C.11: Consistency of the estimates of the relativistic $1s$ energy values, together with their uncertainties for the **hollow nucleus**. The left column is consistency between runs with different basis polynomial orders and the right column is with different z parameter values of the exponential grid. The values on the y-axis are relative to the mean μ and uncertainty σ of the FNC correction (reported below each graph in Hartrees) of all the data points.

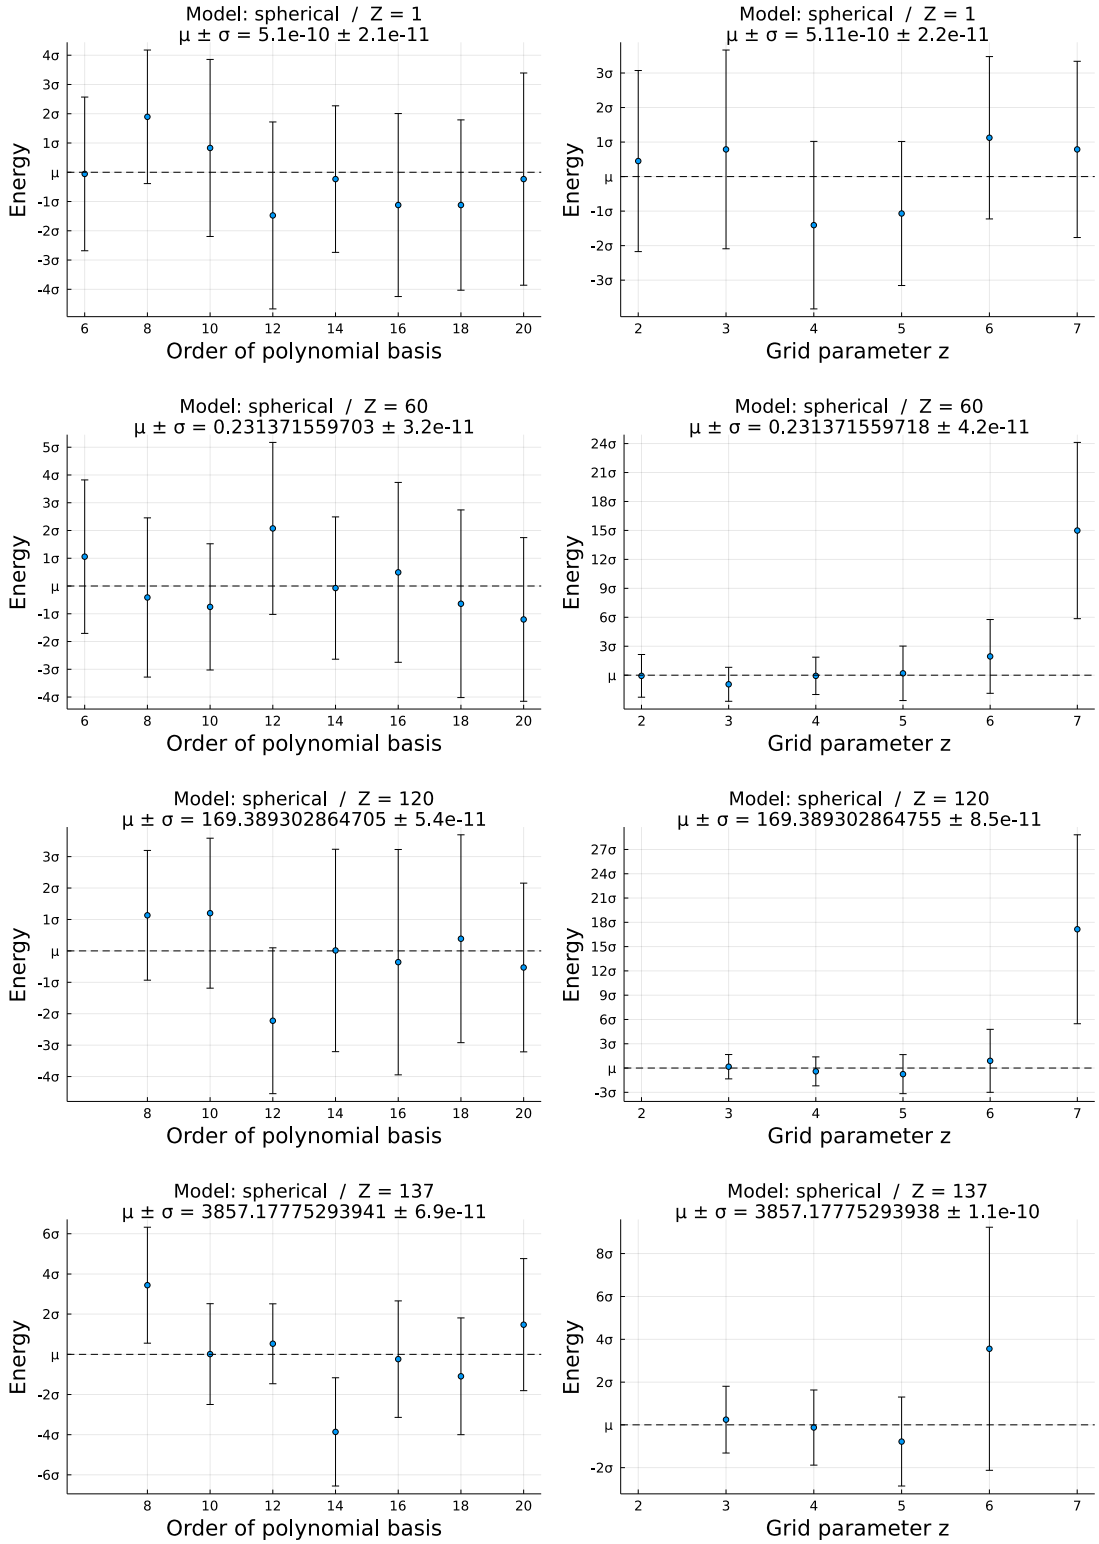

Figure C.12: Consistency of the estimates of the relativistic $1s$ energy values, together with their uncertainties for the **spherical nucleus**. The left column is consistency between runs with different basis polynomial orders and the right column is with different z parameter values of the exponential grid. The values on the y-axis are relative to the mean μ and uncertainty σ of the FNC correction (reported below each graph in Hartrees) of all the data points.

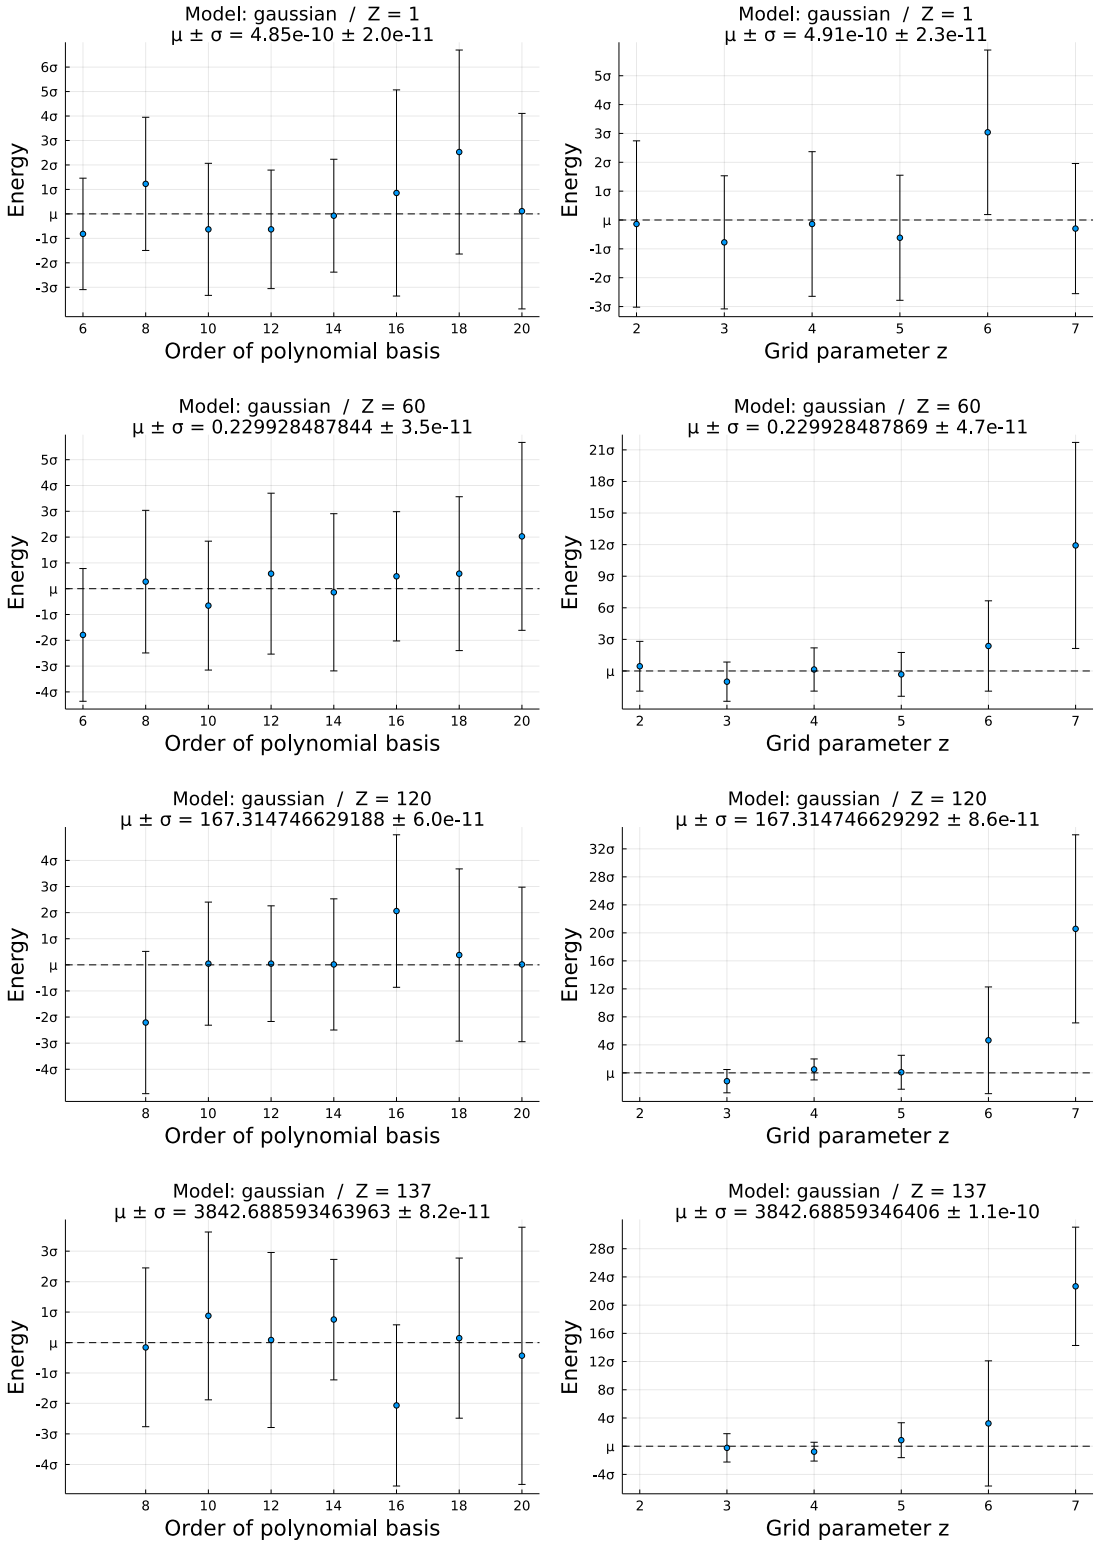

Figure C.13: Consistency of the estimates of the relativistic $1s$ energy values, together with their uncertainties for the **Gaussian nucleus**. The left column is consistency between runs with different basis polynomial orders and the right column is with different z parameter values of the exponential grid. The values on the y-axis are relative to the mean μ and uncertainty σ of the FNC correction (reported below each graph in Hartrees) of all the data points.