

Universality in Quantum Mechanics: A Claim That Cannot Be Tested, and One That Can

Claes Johnson¹

¹KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden ,
claesjohnson@gmail.com

October 14, 2026

The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble. —
P. A. M. Dirac, 1929 [6]

Abstract

Standard quantum mechanics is held universal for chemistry on the basis of a *parameter-free* $3N$ -dimensional Schrödinger equation in configuration space for a system of N electrons, with the nuclear charges as its only input. We ask what grounds that claim, setting it beside RealQM, which replaces that equation by a *new Schrödinger equation in physical three-dimensional space* — the same kinetic and Coulomb energy, a different equation with different solutions. Neither is one equation alone. The standard formalism is Schrödinger plus spin, plus antisymmetry, plus the spin-statistics connection binding them — three components, none of them derived. RealQM adds one, non-overlapping unit electron charge densities, which supplies both the exclusion and the shell capacities 2 and 8. Within chemistry, the domain at issue here, spin proves unnecessary for both of the things it is chiefly invoked to do.

More decisively, the standard claim cannot be tested. In configuration space the cost grows exponentially, so past a few tens of electrons the equation is uncomputable and what is computed instead is a parameterised reduction, whose agreement with measurement does not transfer to the equation behind it. Nor can the approximation hierarchy carry it upward, that hierarchy’s exact limit being the uncomputable object itself. Universality is verified at one to three electrons and assumed beyond, shielded by intractability.

RealQM computes at linear cost, so its universality is empirical: atoms to radon, the periodic table’s row capacities, the X_2 dimers from H_2 to Cs_2 with bond lengths to 7.7%, hydrogen-bond geometries matching measurement, and polypeptides of 153 residues carried with their electrons on a laptop. Its failures are attributable too. The asymmetry is dimensional, not technical, and we close with three tests that would decide.

1 What a universality claim asserts

The claim examined here is not a straw man, and it is worth establishing that before examining it. Dirac stated it in 1929 in the sentence above, and stated it in its strongest form: the laws sufficient for *the whole of chemistry* are completely known. What makes the sentence remarkable is that it asserts the uncomputability in the same breath — the equations are “much too complicated to

be soluble” — and treats that as a difficulty of application rather than as bearing on the claim itself. The programme that followed has largely kept both halves. The term *ab initio* is itself an assertion of derivation from first principles, and the question whether chemistry reduces to the Schrödinger equation remains live rather than settled [23], with Löwdin’s challenge over the $(n+\ell)$ rule presupposing that it does [21, 22].

A careful practitioner will qualify all this, and the qualification should be granted rather than argued with: what is claimed is that the equation is exact *in principle*, not that any particular computation is universal, and the limitations of each method are generally well understood by those using it. That qualification is precisely the subject of this article. An in-principle claim is worth whatever its principle can be made to license, and the question we press is what an in-principle universality licenses when the principle can never be applied to any system for which the claim would be interesting.

What is and is not under examination. One distinction must be fixed at the outset, because the two things are routinely conflated and only one of them is our subject. By *standard quantum mechanics* we mean throughout the N -electron Schrödinger equation itself, posed on $3N$ variables. We do *not* mean the textbook or orbital account of atomic structure — the one-electron hydrogenic states s, p, d, f with capacities $2(2\ell+1)$, filled in the Madelung order — which is an independent-particle approximation laid over the equation, with filling rules laid over that in turn, and which carries no s, p, d, f labels in the equation’s own solutions. The orbital scheme is the easier target and we are not aiming at it.

Nor, it should be said plainly, is the equation’s *correctness* in question here. Where it can be solved it is confirmed to a few parts in 10^5 nonrelativistically, and nothing in this article gives any reason to doubt it. What is examined is the *universality claim* made on its behalf — a claim about scope and about what can be held to account, which is a different kind of assertion from a claim about truth, and which turns out to have a different standing.

A model is universal over a domain of phenomena if three conditions hold.

1. **Species-blindness.** The model makes no reference to the kind of system it is applied to. It does not mention atoms, bonds, metals or proteins; one writes the same thing for helium and for a protein, differing only in Z and the geometry.
2. **Closure under composition.** Joining two systems yields the same model for the union. There is no boundary in the domain at which one must change theories.
3. **Accountability.** The model’s predictions can be compared with measurement across the domain claimed, and the comparison can go against it.

The first two are formal and easily checked; both models satisfy them completely. The third is not formal at all, and it is the one that does the work. A model satisfying only the first two is a *framework*, not a universal theory, because nothing it says about most of its domain can be held to anything.

2 The components of each model

Two models of electronic matter are on the table. Both take the nuclear charge Z as their only supplied number, both use the same law — kinetic energy plus Coulomb interaction — and

neither fits anything. They differ in what must be *added* to that law before it describes matter, and it is worth laying the additions side by side before any question of accuracy arises.

The domain in question throughout this article is **chemistry**: atoms, molecules, materials and biomolecules, their energies and their geometries. Magnetic phenomena — the Zeeman effect, Stern–Gerlach, the electron’s magnetic moment — lie outside it and are not treated here. That restriction is not a convenience: within chemistry all three of the standard formalism’s additions carry weight, spin supplying the capacity two, antisymmetry the shell structure, and the spin–statistics connection the link between them, so the comparison below is made where the components are all doing work.

	standard quantum mechanics	RealQM
state	$\Psi(x_1, \dots, x_N)$ on $\mathbb{R}^{3N}$ [1], complex, antisymmetric	N charge densities on $\mathbb{R}^3$, real, non-overlapping
supplied number	Z	Z
law	kinetic + Coulomb	kinetic + Coulomb
added component 1	spin : a two-valued internal degree of freedom, $s = \frac{1}{2}$ [2, 8]	non-overlap : densities occupy disjoint domains with free interfaces
added component 2	antisymmetry of Ψ , from which Pauli exclusion [3, 4] follows	—
added component 3	spin–statistics [5]: the rule that half-integer spin takes the antisymmetric representation	—
configuration	multiplicity and orbital occupations — an output of full configuration interaction, which <i>cannot be run</i>	the build-up is computable within the model and is carried out to carbon; beyond it the filling is imported, the obstacle being the present outward-only search

Three additions against one. Several remarks are needed to make that count honest in both directions.

Pauli is not a fourth component. Exclusion is a consequence of antisymmetry, not an independent postulate alongside it. What *is* independent is the spin–statistics connection: antisymmetry by itself does not say which particles it applies to. In nonrelativistic quantum mechanics that connection is a postulate; it becomes a theorem only after one adopts relativistic quantum field theory [5], which is a far larger commitment than the equation one began with.

None of the three is derived. Each was added because observation demanded it. The parameter-free Schrödinger equation on its own is content with symmetric solutions and predicts the wrong periodic table; it is the additions that produce chemistry.

Does spin enter as a parameter? This deserves a direct answer, and it has two halves. The *value* does not: $s = \frac{1}{2}$ is a fixed postulated constant, fitted to nothing, and $g = 2$ follows once the Dirac equation is in hand. In that sense spin is structure rather than a parameter. But the *assignment* does, in every calculation actually carried out: one specifies the multiplicity and the orbital occupations — that this system is a singlet, that one a triplet — and the computation does not determine them. In principle a full configuration interaction in a complete basis would return the spin state as an output rather than take it as an input, so the assignment is not a parameter of the theory. That defence, however, rests on exactly the calculation which cannot be performed, which is the subject of Section 4.

Both import a configuration, and the imports are not alike. RealQM derives the shell capacities at every Z and performs the build-up explicitly through lithium, beryllium, boron and carbon; beyond carbon the filling is imported. So a configuration is supplied in each case, and it would be easy to call that a draw. It is not one, and the difference is the subject of this article.

The standard formalism’s multiplicity is an output of a full configuration interaction which *cannot be run*, now or by any method expected to exist. The limitation is permanent. RealQM’s build-up, by contrast, is an ordinary computation within the model — it is what produces the configurations through carbon — and what stops it beyond carbon is that the procedure presently used adds charge outward, so it cannot place a charge inside a shell already occupied. That is a restriction on the search, not on the model, and it is removed by writing a different search rather than by waiting for an intractable problem to become tractable.

So the honest statement is that the Aufbau build-up is computable in RealQM in principle and carried out to carbon in practice. One of the two limitations has a known route out of it; the other does not.

What the single postulate buys. Non-overlap does the work of components 2 and 3 together: electrons exclude one another because they cannot occupy the same region, so no antisymmetry need be imposed and no statistics rule is needed to say whom it applies to. It also supplies the shell capacities, which the standard formalism obtains from spin multiplied by orbital degeneracy — the antipodal pair gives 2, and the octet $8 = 4+4$ follows from the same packing rule, two dual tetrahedra — the decomposition of Linnett’s double quartet [19], reached here from Coulomb interaction and non-overlap rather than from spin — since a lone tetrahedron has no inversion centre and cannot close.

Is spin necessary? The question has no single answer because it is three questions.

- **For exclusion: no.** Territory does the work, and this is not a conjecture — it is what every atomic result in Section 5 is computed from.
- **For the capacity two: no.** The antipodal pair gives it geometrically, and the same rule gives eight. Neither borrows the spin degeneracy.
- **For the two-valued magnetic response: outside this article.** A scalar charge density gives one undeflected Stern–Gerlach spot, not two with no middle, so something further is required there. Whether that something must be spin as the standard formalism has it is a real question and not one we settle; it belongs to the magnetic phenomena excluded above.

Within chemistry, then, spin is not necessary for either of the things it is chiefly invoked to do. And it should be said that neither model *explains* spin: the standard formalism posits it. Its presence there cannot be defended as something nature requires and an alternative merely lacks — it is an addition, like the other two, made because observation demanded it. What is distinctive is that the standard formalism needs it twice over: once as spin, and again as the statistics rule which spin is then made to govern.

A universal model should cover more with less. On that measure the comparison above favours the one-postulate model before any question of accuracy is reached, and the rest of this article asks whether either model’s universality can be tested at all.

3 How many parameters, and where from

Being wholly free of parameters is an ideal rather than a practice. Every model that reaches a large system does so through a reduction, and reductions carry constants. So the useful questions are not whether a model has parameters but *how many, to reach what scale*, and — decisively — *where they come from*.

That last question admits two answers which are not epistemically alike.

- **Upscaled:** the constant is computed from the model itself at a finer scale. The coarse model then has a parameter, but it was derived rather than fitted, and a prediction made with it remains a prediction. This is ordinary multiscale practice and it is legitimate.
- **Fitted to observation:** the constant is adjusted so that the model reproduces measured data. If the data fitted is the data later predicted, the agreement is circular and tests nothing. If it is a *different* observable, the agreement is a weaker but real test: one quantity is spent to buy another.

Both models use both kinds. Setting them out by scale is more informative than any global claim about parameter-freedom:

scale	RealQM	standard quantum mechanics
atoms	none ; Z alone, all-electron on the bare kernel $-Z/r$	none fitted to experiment — but basis-set exponents are optimised by minimising <i>atomic</i> energies, so an atomic energy computed in that basis is not an independent prediction
small molecules	none , where the atoms are treated all-electron	none fitted to experiment in the CCSD(T) near-basis-set-limit window; the basis is upscaled from atomic calculations, though its form and truncation are chosen and the basis-set limit itself is reached by a fitted extrapolation
larger molecules, materials	none in the model , which is defined all-electron on the bare kernel. In practice a core radius r_c avoids resolving the $1s$ on a molecular mesh — a numerical device, like the mesh itself, and not a component. But it is not demonstrably benign: see below	a choice of exchange–correlation functional, several of them fitted to benchmark sets of the properties later computed
dispersion	absent from the model	an added term fitted to reference data
biomolecules	r_c , plus the secondary structure (which residues are helical, driving a hydrogen-bond bias) and a surface-area term standing in for the hydrophobic effect. Tertiary packing is <i>not</i> biased	not reachable; classical force fields reach it, wholly fitted to experiment and to quantum calculation

Three things follow, and the second is not favourable to the argument of this article.

r_c is a **numerical device, not a component** — but it is not a clean one. The core radius deserves separating from the other entries, because it is a different kind of object. The model needs no r_c : every atom in Section 5 is computed all-electron on the bare kernel $-Z/r$, hydrogen to krypton, which settles that r_c is not part of the formulation. What it does is numerical. A molecule needs a box spanning the whole molecule, while the innermost shell contracts as $r_{1s} \approx 4.3/Z$, and a single uniform mesh cannot do both; replacing the resolved core by a reduced kernel and an exclusion radius buys the molecular calculation at the cost of a length. In that respect r_c belongs with the mesh spacing rather than with a fitted constant, and a model carrying only numerical parameters carries none that matter for the present argument.

That is the favourable reading and we should not take it without the qualification. A benign numerical parameter does not change answers qualitatively, and this one does. On the hydroxyl fragment, with the oxygen valence split and the boundaries free, $r_c = 0$ gives a dipole of 1.98 D

with the oxygen negative, against a measured 1.66 D; at $r_c = 0.5$ the same configuration gives 0.34 D *with the polarity inverted*. A mesh spacing that reversed the sign of a dipole would be called a broken discretisation, and the same standard applies here.

It is worth being quantitative about what r_c actually buys, because the usual phrase — that all-electron molecules are too expensive — turns out to name the wrong cost. For a protein built from hydrogen, carbon, nitrogen and oxygen the *electron count* barely moves: per residue the reduced description carries about 35 electrons against 59 all-electron, and for a 153-residue chain 5340 against 9042, a factor 1.7. That is not what makes the calculation hard. The cost is the *mesh*. The run above uses 200^3 cells on a 100 bohr box, so $h = 0.5$ bohr, while resolving an oxygen $1s$ — $r_{1s} \approx 4.3/Z \approx 0.54$ bohr, spanned by some ten cells — requires $h \approx 0.054$. That is a factor 9.3 in each direction and therefore some 800 in cells, a 200^3 grid against roughly 1860^3 ; with cost going as electrons times mesh points the product is of order 10^3 .

But that factor is only needed within about a bohr of each heavy nucleus. With roughly twelve hundred carbon, nitrogen and oxygen atoms, core-refined patches of radius one bohr at $h = 0.054$ come to some 6×10^7 cells in total, against 6.4×10^9 for refining everywhere, which places an all-electron calculation of the same chain at something like fifteen times the present cost rather than a thousand. So r_c is not required by the physics, as the all-electron atomic results show, and it is not required by the cost either. It is required by the use of a *uniform* grid. On a mesh with local refinement it would disappear rather than be calibrated — which, given that it reverses a dipole, is the better outcome. Two details make the present situation less comfortable still: at $h = 0.5$ the oxygen core radius of 0.8 bohr spans one and a half cells, so the exclusion sphere is represented by a cell or two and the polarity inversion may itself be a resolution artefact; and the $r_c \rightarrow 0$ limit has never been taken for a molecule, that being precisely the calculation r_c exists to avoid.

So the honest position is that RealQM has no model parameters at any scale, and one numerical device which is an artefact of uniform meshing and whose convergence has not been demonstrated where it is used.

What “nothing fitted” does and does not mean for the standard formalism. It means nothing fitted to experiment, and that is worth less than it sounds. Three of the four qualifying remarks are mechanical: the basis form is a choice rather than a derivation, Gaussians with a chosen primitive count and contraction pattern; the basis-set limit, which is the rung supposed to be free of constants, is approached through an extrapolation formula carrying fitted ones; and effective core potentials for the heavier elements are fitted to reference data outright. The fourth is more serious. Basis-set exponents are determined by minimising *atomic* energies, so an atomic energy computed in such a basis is not an independent prediction — the variational principle largely guarantees a good number from a basis tuned to that very system. The circularity is milder for molecules, which is why basis-set superposition error exists as a named effect; but it means the atomic scale, where the standard formalism looks cleanest, is in fact its least independent test.

The clean regimes coincide, and neither model is tested outside them by its own successes. For atoms and small molecules both models can be run with no constant fitted to the predicted quantity, and both should be believed there, with the reservation just entered. RealQM’s advantage in that band is coverage rather than kind: it runs all-electron across the table and into the fourth row at a few per cent, where the standard formalism’s clean window

is a few tens of atoms at a far higher precision. That is a genuine trade and not a defeat either way.

RealQM’s parameters become circular at exactly the scale where its scope claim is most striking. The core radius is numerical, as just discussed. At the protein scale the secondary structure is supplied — which segments are helical, which then drives a hydrogen-bond bias — together with a two-parameter surface term standing in for the hydrophobic effect. The tertiary packing is not biased and is computed, so what those runs cannot be offered as is evidence that the model predicts secondary structure. The honest reading is that RealQM’s parameter count rises from zero to one to three as the scale grows, and that the last step crosses from upscaled to fitted. Any claim resting on the largest systems must be read in that light.

But the crossing happens earlier for the standard formalism, and for everything. Outside the CCSD(T) window — which is to say for every system in materials science, catalysis and biology — the computed object is a functional-bearing reduction, and the functionals are in substantial part fitted to the properties they are then used to compute. The dispersion term is fitted outright. So on the parameter question the comparison is not parameter-free against parameterised; it is *one parameter sourced from a different observable, out to materials and solids*, against *a fitted functional from small molecules onward*. Where each model stops being testable by its own output is the real content of the comparison, and RealQM’s crossing comes later.

4 Why the universality of standard quantum mechanics cannot be tested

4.1 The equation is uncomputable where the claim is interesting

The state is a function on $\mathbb{R}^{3N}$, and representing it costs exponentially in N : full configuration interaction in a basis of M functions scales as $\binom{M}{N}^2$ in the worst case, and the exact solution is in practice out of reach past roughly twenty electrons. This is not a matter of waiting for better machines. The ground-state problem for a local Hamiltonian is QMA-complete [10, 9], so no efficient general algorithm is believed to exist even on a quantum computer; the barrier has the character of a theorem rather than of an engineering limit.

The systems for which universality is an interesting claim — a catalyst, a protein, a crystal, a cell — carry 10^3 to 10^{23} electrons. For every one of them the model cannot be solved.

4.2 What is computed instead is not the model

Since the equation cannot be solved, it is not the equation that is compared with experiment. What is compared is a reduction, and every reduction in use carries parameters. Gaussian basis sets are built by optimising exponents and contraction coefficients atom by atom, and the complete-basis-set limit is reached by extrapolation through a fitted functional form. Coupled-cluster theory [13] truncates the excitation operator, and CCSD(T)’s perturbative triples correction [14] is not variational and carries no bound — its celebrated kcal/mol accuracy is an empirical regularity over well-behaved closed-shell molecules which degrades without warning for multireference systems. Density functional theory [11, 12], which performs essentially all practical calculation on large systems, selects an exchange–correlation functional from a zoo in

which B3LYP, PBE, M06 and ω B97X are different theories with different failures, several fitted to benchmark sets outright.

We should be fair about the strongest case. There is a window — small closed-shell molecules of up to a few tens of atoms — in which CCSD(T) near the basis-set limit is nearly parameter-free and agrees with measurement to about a kcal/mol across a wide range of chemistry. That is a real achievement and we do not dispute it. But it is a window, and outside it, which is to say for almost everything, the computed object is a functional-bearing reduction.

4.3 The hierarchy cannot transfer verification upward

It is often said that this does not matter, because the approximations form a convergent hierarchy — Hartree–Fock, MP2, CCSD, CCSD(T), full CI — whose limit is the exact solution, so the hierarchy connects what is computed to what is claimed.

It does not, for a structural reason. The hierarchy has two ends and both are unavailable. Its exact limit is the uncomputable object itself, so one can never stand at the top and check. Its lower rungs, which are what gets computed, each carry parameters, so agreement there is agreement of a parameterised procedure. There is no rung which is both reachable and parameter-free, and consequently no step at which empirical success attaches to the parameter-free equation rather than to the reduction. The existence of a convergent sequence establishes that the reductions approach something; it does not establish that the something is nature, since that identification is exactly what was to be tested.

4.4 The practice this produces

The consequence is visible in how discrepancies are handled. When density functional theory fails on dispersion — and semilocal and hybrid functionals do fail on it, systematically — the response is to add an empirical correction, a fitted term [15], and the failure is charged to the functional. When a coupled-cluster result is wrong for a transition-metal complex it is charged to the truncation or the basis. In no case is it charged to the Schrödinger equation, and this is not dogmatism: it is forced, because no computation of the equation exists against which the equation could be found wanting.

A claim that no measurement can contradict is not thereby strong; it has been removed from the class of claims that measurement bears on. Over the great majority of its asserted domain that is the position of the universality of standard quantum mechanics — verified superbly at one to three electrons, assumed beyond.

5 Why the universality of RealQM can be tested

RealQM is in the opposite position on the one point that matters, for one reason: it is computable where the claim is interesting. The state is N densities on $\mathbb{R}^3$, so the cost is the electron number times the number of mesh points — linear in N , not exponential. A molecule, a cluster, a crystal, a protein is the same computation with more kernels and more domains. The model itself, with Z as its only input and no reduction interposed, can be run on the systems in question and the output compared with measurement.

That makes universality an empirical question in the sense of Section 1(3). What follows is the record. It is organised by scale, because the point is not any single number but that the same four-line scheme, unchanged, reaches every scale at which chemistry happens.

5.1 Atoms

Hydrogen is an identity, not a test. At $N = 1$ the configuration space *is* physical space, there is no exclusion and no partition, and the two equations coincide exactly — the only electron number for which they do. Both give $E = -0.5$ Ha and $\alpha = 4.5a_0^3$. We record it because it locates where the two models can first differ — at two electrons, where the exclusion postulate enters.

Helium. The first real test — the first electron number at which the two models can differ at all. At two electrons the exclusion postulate can be imposed in two ways, and both are worth reporting, because they fail in opposite directions and together they bound what the construction can give.

	E (Ha)	error
RealQM, one shared density, pair counted once	−2.8611	+1.5%
Hartree–Fock	−2.8617	+1.5%
nonrelativistic benchmark [16]	−2.903724	—
RealQM, two non-overlapping domains	−3.0265	−4.2%

The benchmark lies between the model’s own two constructions. Letting the two electrons share one density, with their single pair interaction counted once, under-binds by 1.5%; confining them to disjoint half-spaces over-binds by 4.2%. A shared density gives no mutual avoidance beyond the mean field, disjoint domains give complete and permanent avoidance, and the true state avoids partially — so non-overlap behaves as a dial between no avoidance and total avoidance, with helium’s energy inside the span. Only the second construction generalises: beyond two electrons a single shared density has no shell structure, so no atom past helium and no periodic table, and restoring distinct states requires orthogonality, which is antisymmetry again. The −4.2% is therefore the figure that characterises the model, and the He-like series behaves the same way — Li^+ comes out −7.5031 against −7.2799, over-bound by 3.1%.

Two earlier claims withdrawn. Previous versions of this section reported −2.91 at 0.24% and concluded that territorial exclusion beats the mean-field theory beside it, recovering about two thirds of the correlation energy. Neither statement survives.

The first was a relaxation transient. The energy in that computation rises monotonically and passes *through* −2.9037 on its way to the construction’s own minimum; driven to its stopping point the same program gives −2.826 and is still rising, and its kernel softening is some 3.5× softer than the value calibrated against the exact one-electron energy. A value the iteration transits is not a result.

The second cannot be rescued even in principle, because the shared-density construction above *is* the Hartree functional, and for a closed two-electron shell Hartree coincides with Hartree–Fock. Its computed −2.8611 reproduces the Hartree–Fock −2.861680 to 0.02%, which is how we now check the solver. So RealQM does not beat Hartree–Fock on helium: one end of its own bracket essentially *is* Hartree–Fock, and the other end overshoots. No correlation energy is recovered by exclusion; what exclusion does is replace a mean field by a geometry, and on helium’s total energy the two land on opposite sides at comparable distance.

What the benchmark is, and why it matters here. The -2.903724 above is not a measurement. It is the computed solution of the nonrelativistic Schrödinger equation at infinite nuclear mass — the equation under test, solved by a specially constructed expansion. The measured quantity is the pair of ionization energies, and summing NIST’s two values gives -2.903386 Ha; the remaining 0.00034 Ha is *removed by calculation*, being the relativistic, QED and nuclear-motion content the measurement has and the benchmark must not. Benchmark and measurement part company in the fourth decimal, and the gap is supplied by theory.

The same holds, less sharply, for the atomic table below. Its reference energies combine experimental data with *ab initio* results — the correlation contribution obtained by comparing experiment against relativistic calculations — and past argon they are computed outright at Hartree–Fock level [17]. A nonrelativistic stationary-nucleus total electronic energy is not an observable; no instrument reads one. So the per-cent figures here measure distance from a semi-empirical composite that already contains standard quantum mechanics, not distance from nature, and they cannot be read as a verdict of one theory against experiment. This does not make them worthless — the composites are accurate and agreement to a per cent against them is a real constraint — but it sets what they mean, and it is why the observables we rest the universality claim on are of a different kind: ionization energies, geometries, shell radii, and ordering.

The table. Computed all-electron on the bare kernel $-Z/r$, with no core radius, no fitted length and nothing adjusted per element: through argon the mean relative error is 1.6% with maximum 4.5%, and 13 of 17 atoms fall within 2.5%. Reference energies throughout are the near-exact nonrelativistic values of Chakravorty *et al.* [17]. Extended to the fourth row on a mesh chosen by Z , the mean is 3.3%, every entry under-bound in the direction the spherical reduction predicts. Energies are available through to radon.

The periodic table’s capacities. Two filled packings appear and no third: the antipodal pair, 2, and the octet, $8 = 4+4$, two dual tetrahedra. A lone tetrahedron has no inversion centre and so cannot close a shell; its dual completes the cube, which is centrosymmetric and admits no further dual. Row capacities are sums of these two alone,

$$2, \quad 8, \quad 18 = 2+8+8, \quad 32 = 8+8+8+8,$$

and a period ends when its capacity is exhausted, giving the closures 2, 10, 18, 36, 54, 86, 118. The capacities obey a parity rule — remainder modulo 8 is only ever 0 or 2, never 4 or 6, so each is a whole number of octets carrying at most one pair, the pair appearing exactly when the shell index is odd — and admit a closed form, $\text{cap}(n) = 8\lfloor n/2 \rfloor \lceil n/2 \rceil + 2(n \bmod 2)$, which equals $2n^2$ for every n without ever invoking it. The account coincides with the hydrogenic degeneracy rather than borrowing it.

Spectra and valence quantities. Excited-state spectra are computed for the alkalis, the alkaline earths and helium, including the helium ground-state-to-orthohelium excitation, and first ionization energies for the first ten elements, correct in sign throughout and to a mean 18% in magnitude (Section 5.5). Electronegativity comes out natively: Allen’s configuration energy [20] reproduces the ordering across period 2 exactly.

5.2 Molecules

The covalent bond, as one member of a family. H₂ is usually reported on its own, which invites a flattering comparison. We report it instead as the $R = 0$ member of the single-valence dimer series X₂, since a +1 kernel with no core *is* X with ionic radius zero, and the alkalis are the same construction with a core of radius R excluded. Within each row, R is fixed by requiring the pseudo-atom to reproduce that element’s *measured* ionization energy; the bond length and the binding energy are then predictions, and the measured ionic radius of M⁺ is an independent check on R that nothing was fitted to. Computed on the axisymmetric solver, which is an exact reduction for two domains on an axis:

	R	$r_{\text{ion}}(\text{M}^+)$	IE	IE obs	D_{min}	R_e obs	E_{bind} / obs
H ₂	0	0	13.61	13.60	1.350	1.401	+58%
Li ₂	2.00	1.44	5.42	5.39	5.425	5.051	+36%
Na ₂	2.16	1.93	5.11	5.14	5.425	5.818	+97%
K ₂	2.79	2.61	4.35	4.34	6.825	7.416	+96%
Rb ₂	2.97	2.87	4.17	4.18	7.175	7.955	+101%
Cs ₂	3.30	3.16	3.89	3.89	7.875	8.783	+95%

(lengths in a_0 , energies in eV.) Read as a family, the result is one statement rather than six: **bond lengths come out to 7.7% on average, with the correct trend down the group, and binding energies are over by about a factor of two throughout.** The two errors differ so much in size because the binding energy is a small difference of large totals — for H₂ it is a seventh of the total, so a few per cent on the total necessarily becomes tens of per cent on the difference — and H₂ is not special: it carries the family’s error, not a better one.

The independent check is worth isolating. The core radius needed to reproduce each measured ionization energy lands close to the measured ionic radius of the ion that remains, the gap falling from $0.56 a_0$ at lithium to $0.10\text{--}0.14 a_0$ for Rb and Cs. Nothing in the calculation was told what an ionic radius is.

Two limitations belong with the table. The five alkali rows share one mesh whose admissible separations are spaced $0.35 a_0$ apart, and Li₂ and Na₂ both land on $D = 5.425$ although their measured bond lengths differ by $0.77 a_0$: *adjacent species are not resolved*, and the resulting Na₂ > Li₂ binding order is an artefact of that spacing rather than a prediction. And of the H₂ curve only R_e and D_e are measured quantities at all — the energy at arbitrary R is itself computed, so the textbook six-decimal agreement for H₂ belongs to a computation and not to a measurement.

A previous version of this section reported H₂ alone at 94% of the measured dissociation energy. That figure is withdrawn: it evaluated the model at $R = 2.0 a_0$ rather than at equilibrium while comparing against the exact energy at $R = 1.401$, it compared an electronic binding energy against D_0 rather than D_e , and its dissociated limit was itself 3.8% too low, so that part of the error cancelled in the subtraction.

Water. The dipole moment and the bond angle are the two quantities that test whether the charge is in the right place. With the oxygen valence split into angular sectors, one toward each hydrogen and one pointing away:

quantity	RealQM	experiment	ratio
dipole moment μ	1.836 D	1.85 D	99%
H–O–H angle	101.6°	104.5°	97%
equilibrium O–H bond	1.814 a_0	1.814 a_0	exact

The lone pair is not put in by hand; the sector pointing away from both hydrogens is where the energy puts it. The equilibrium bond length comes out on the measured value.

5.3 Condensed matter and reactions

The hydrogen bond is where molecular matter becomes material, and it is geometric:

system	RealQM	reference
water dimer	H $\cdots$ O = 1.87, O–O = 2.95 Å	matches measurement
formamide dimer	N–H = 1.01, H $\cdots$ O = 2.04, N $\cdots$ O = 2.97 Å	matches measurement
ice I _c , 216 molecules	O–O = 2.73 Å	2.76 Å

Beyond hydrogen bonding, the same scheme gives metallic bonding — a 100-atom lithium lattice, perturbed, restores itself — ionic crystals and their melting in sodium chloride, conduction barriers for distinct arrangements of charge, and chemical reaction paths: $\text{H} + \text{H}_2 \rightarrow \text{H}_2 + \text{H}$ proceeds by bond exchange through the H_3 transition state, with proton transfer, ion formation and salt dissolution computed on the same footing.

5.4 Biomolecules

Nothing in the model changes at this scale; only the number of domains. One qualification must precede the table, because without it the entries would be read for more than they show. At this scale the runs are not blind. Two things are supplied: the *secondary structure*, i.e. which residues are helical, from which an $i \rightarrow i+4$ hydrogen-bond bias is generated; and a surface-area term with two constants standing in for the hydrophobic effect. The tertiary packing — how the helices arrange against one another — is *not* biased and is computed. Both supplied terms stand in for identifiable missing physics. Blind assembly fails because electron repulsion dominates at intermediate separations of three to eight Ångström while the attractive hydrogen bond reaches only about two, so nothing nucleates; and the hydrophobic effect originates in water entropy, to which a ground-state model of electron density has no access. What these entries therefore show is that the quantum model produces *correct hydrogen-bond geometry, and tertiary arrangement, at the scale of a protein* once the secondary structure is given — not that it predicts secondary structure or folds from sequence. The parameter-free hydrogen-bond results are the dimers and the ice of the previous section; these are the same physics carried to a larger system with the approach supplied.

system	result	reference
DNA G:C base pair	2 of 3 Watson–Crick bonds at ~ 2.1 Å	1.9 Å
RNA hairpin, 12 nt	all 4 stem pairs form, G1:C12 = 1.99 Å	2.0 Å
Trp-cage, 20 residues	4 of 5 helix H-bonds, $R_g = 7.2$ Å	near native
Villin HP35, 35 residues	3 helices, H-bonds at 2.0 Å, core F10–F17 = 6.1 Å	3-helix bundle
Myoglobin, 153 residues	helices A = 1.51, G = 2.15, H = 2.68 Å	partial fold
coiled coil, 2×15 residues	helix separation $12 \rightarrow 7.5$ Å, blind	water-driven docking

The RNA hairpin assembles its own secondary structure from a U-shaped start: stem paired, loop open, correct topology. The coiled-coil docking is blind — no bias toward the bound state — and is driven by the water.

5.5 Attribution: a prediction made, tested, and refuted

Accountability has a second half that is easy to overlook. A model is held to account not only when its predictions are checked but when its *failures* can be located — assigned to the model, or to a reduction of it, with the assignment itself open to test. The standard formalism cannot do this, as Section 4 argued: a discrepancy is charged to the functional or the basis because no computation of the equation exists to charge it to. What follows is the opposite case, carried through to its conclusion, and the conclusion went against us.

The discrepancy, and a bracket. Ionization energies come out badly *in the radial solver*. Beyond lithium the valence charge is too weakly bound, and by $Z \approx 14$ it is unbound altogether; the material was withdrawn from the companion article on atoms for that reason. Whether this indicts the model or a reduction of it is decided by a measurement rather than by argument — and, as the next paragraph shows, the reduction is where it belongs. The radial solver has no angular coordinate, so it cannot represent the octet as the model describes it — 4+4 at one radius, two dual tetrahedra angularly separated — and can represent only two alternatives. Computed for neon by the same Δ SCF route throughout, so that the comparison is like for like:

$$\underbrace{13.33 \text{ eV}}_{\text{radial } 4+4, 2+4+4 \rightarrow 2+4+3} < \underbrace{21.56 \text{ eV}}_{\text{measured}} < \underbrace{34.48 \text{ eV}}_{\text{single 8-shell, } 2+8 \rightarrow 2+7}$$

The measured value lies inside, a factor 1.6 either side, and the same bracketing holds for every noble gas from neon to radon. Two cross-checks confirm the solver rather than the arithmetic: the radial 4+4 total of -131.83 Ha reproduces the -131.731 tabulated in the companion article on atoms [28], and the single eight-shell over-binds by 11%, inside the 9–22% range recorded for compacting an octet onto one shell.

The same difference with the valence free in three dimensions. The radial reduction is not the model. Give the valence angular freedom on a three-dimensional grid over homogenised inner shells — with the configurations the model itself supplies and the same Δ SCF difference throughout — and the first ionization energies of the first ten elements come out with the correct sign everywhere and a mean absolute error of 18% (Hartree):

	neutral $\rightarrow$ cation	3D Δ SCF	measured	ratio
He	2 $\rightarrow$ 1	0.90	0.9036	1.00
Li	2+1 $\rightarrow$ 2	0.198	0.1981	1.00
Be	2+2 $\rightarrow$ 2+1	0.43	0.3425	1.26
B	2+2+1 $\rightarrow$ 2+2	0.20	0.3049	0.66
C	2+4 $\rightarrow$ 2+2+1	0.40	0.4138	0.97
N	2+2+3 $\rightarrow$ 2+4	0.67	0.5341	1.25
O	2+4+2 $\rightarrow$ 2+4+1	0.53	0.5001	1.06
F	2+4+3 $\rightarrow$ 2+4+2	0.84	0.6403	1.31
Ne	2+4+4 $\rightarrow$ 2+4+3	1.20	0.7925	1.51
Na	2+4+4+1 $\rightarrow$ 2+4+4	0.19	0.1886	1.01

Lithium, sodium, carbon, oxygen and helium land within a few per cent; the error grows across the row and is worst at fluorine and neon, where the model over-binds the valence by half. That is a graded, locatable failure rather than a collapse, and it places the radial solver’s unbinding where it belongs: in the loss of the *valence’s* angular freedom. The core may be homogenised with little cost — it must be, since a uniform grid cannot resolve a $1s$ shell at a_0/Z and a valence shell at 8 a.u. simultaneously — but the outermost shell may not. A second construction avoids the difference altogether. Screen the core to its net charge, give the valence its own domain outside a radius R , and the ionization energy is $E_{\text{val}}(n) - E_{\text{val}}(n-1)$ with the core energy cancelling identically rather than numerically; at $R = 0$ and kernel +2 this returns helium’s 0.903 against a measured 0.9036, with nothing supplied. Its limitation is the converse of a parameter: R must be given, and there is no quantity inside the radial model that supplies it. The core–valence interface of a space-filling radial partition is not the screening radius — the first is a property of homogenising each shell onto a sphere, the second is where the enclosed charge leaves the valence seeing net +1 — and the two do not track. Taken from the interface, R gives lithium 0.1968 against 0.1981 but sodium only half its value, the interface having grown to 5.38 a.u. where the screening radius must be near 2.1. The interfaces grow by roughly 2 a.u. per shell (1.90, 4.01, 5.38, 9.18, 9.42 for one through five core shells) where physical core radii saturate near 2–3: a partition that gives every shell a nearly fixed radial width, rather than contracting it with the charge that shell sees, misplaces every interface past the first. So R is properly fitted in this construction, and the parameter-free route to an ionization energy is not to predict R but to dispense with the pseudo-kernel altogether, as the table above does.

What a fitted length can still be held to. Not every fitted parameter carries the same weight, and the distinction is worth making precisely because this article turns on it. A parameter is uninformative when it is fitted to the same observable it is then said to predict: five radii chosen to reproduce five ionization energies assert nothing. The construction above can be put in a stronger position. Remove the outer cutoff — solved on $R < r < \infty$ the problem is well posed and R is the *only* length in it — and ask not for a radius per atom but for one constant across a column. Against Shannon’s six-coordinate crystal radii [26], the R that reproduces each measured first ionization energy is a fixed fraction of the ion’s measured radius:

$$R = 0.439 r_{\text{ion}} \implies \text{Li } 1.04, \quad \text{Na } 0.97, \quad \text{K } 1.01, \quad \text{Rb } 1.00, \quad \text{Cs } 1.02$$

as ratios of computed to measured, all five alkalis within 4% and four within 2% from a single number. On Shannon’s effective scale the constant becomes 0.484 and the agreement is comparable, so the result does not depend on the radius convention — only the constant moves with

it. One parameter standing against five observables is a claim that can fail, and it is the form in which this construction is worth reporting: the model supplies the function $IE(R)$, and given a measured ionic radius that function returns the ionization energy. What it does not do is produce the radius, so this is a correlation between two measured quantities mediated by the model rather than a prediction from Z alone — which is exactly why the all-electron route of the preceding table, needing no R at all, carries the parameter-free claim and this one does not.

The prediction. The physical arrangement is geometrically intermediate between the two the reduction admits, so a three-dimensional angular octet should land inside the bracket, near 21.56 eV. That is a falsifiable prediction about where the failure belongs, and it is the kind of statement the model of Section 2 can be made to answer.

The result. We ran it: eight co-located unit domains at cube-corner directions on a pseudo-kernel of +8, with the cation realised as the same eight sectors with one emptied, so that kernel and partition are identical between the two solves and the core’s self-energy cancels by Gauss. Across core radii $r_c = 0.30, 0.40$ and 0.50 a.u. the ionization energy came out at 107.9, 96.9 and 89.0 eV — four to five times the measured value, and *above the entire bracket* rather than inside it. Nor can a different r_c rescue it: the steps are -10.9 then -8.0 eV for equal increments, a decelerating trend whose ratio extrapolates to an asymptote near 67 eV. The over-binding shows in the totals, the valence sector coming out at $-55.7, -47.5$ and -41.8 Ha where it should be about -35 .

So the attribution fails, and one escape route closes with it. The bracket survives as a measured fact; the inference drawn from it does not. It might be supposed that the three configurations involved were simply chosen badly — and at neon they are supplied, by the balanced filling rule rather than derived. But the octet construction’s failure is not a configuration artefact: the +8 pseudo-kernel over-binds the valence sector to $-55.7, -47.5$ and -41.8 Ha where it should be about -35 , at every r_c tried. What the three-dimensional all-electron result above does show is that the *radial* solver’s unbinding is a property of that reduction: carbon, the last element whose configuration the model’s own build-up *derives*, comes out at $0.26\times$ the measured value radially and $0.97\times$ in three dimensions, with the same ΔSCF difference and the same derived configuration in both.

What this costs, and what it shows. Two attributions were at issue and they come apart. The specific one — that a three-dimensional angular octet would land inside the bracket — is refuted, and no reading of the result rescues it. The general one, that a failure of this model can be charged to a reduction rather than to the model, survives and is now demonstrated rather than asserted: the radial solver’s unbinding is the reduction’s, since the same difference computed in three dimensions with the same derived configuration recovers carbon to 3%. What remains genuinely open is narrower and sharper than either. At carbon the build-up derives $2+2+2$ while the energy functional, minimised over configurations, prefers a single six-shell 21.77 Ha lower and 56% wrong. An ionization energy is a difference of two quantities that each ought to be a minimum over configurations; if the minimisation points elsewhere, no difference built from that functional can be trusted. Total energies survive precisely because the configuration is supplied — they are evaluations at a given point, not minimisations — so good totals and bad differences are one defect seen from two sides.

We report this at length, rather than removing it, because it is what the argument of this article looks like when it runs. A prediction was stated, a reference value not in dispute was chosen, the calculation was performed, and the model lost. The universality claim examined in Section 4 cannot be placed in that position at all — not because it is better supported, but because there is nothing to run.

5.6 The scope is exercised, not asserted

This is the comparison that bears directly on universality, because it is about what each method can be *held to*:

method	largest system carried	hardware	time
Gaussian / VASP (DFT, HF) [11, 12]	cannot fold proteins	HPC cluster	days per point
GROMACS / OpenMM (MD)	20 residues	GPU clusters	weeks to months
Anton 2 (MD) [24]	80 residues	\$100M custom ASIC	months per protein
RealQM	153 residues, all electrons	laptop GPU	minutes to hours

The quantum-chemical methods — the ones carrying the universality claim — appear in the first row, and they do not reach the problem. The two middle rows reach it with empirical force fields fitted to experiment and quantum calculation, so they are not tests of quantum mechanics at all. Machine-learning structure prediction [25] is a different kind of thing again: it predicts a structure from sequence without computing the physics, so it is not a test of any model of matter. RealQM carries the system from the charges and Z on commodity hardware — with, at this scale, the contact bias noted above. The claim here is about *scale*, which is what accountability requires: a model that cannot be evaluated on a protein cannot be checked against one. Whatever one concludes about accuracy, RealQM’s scope is exercised rather than asserted, and that is the condition Section 1(3) asks for.

6 Scope and fidelity

One bound must be stated, because it marks the limit of what the record above establishes. RealQM is a fixed ansatz. Refining its mesh converges to the ansatz’s own answer, not toward nature’s, and there is no next order to add — no analogue of going from CCSD to CCSD(T) to full CI. So the claim supported here is universality of **scope**: the model applies across every scale of chemistry, is checked there, and agrees to a few per cent on energies and considerably better on geometry. It is not a claim of universal **fidelity**, which would require a convergence ladder that the model does not have. Being a testable claim, it is also refutable, and in places it has been refuted; that is the price and the point of being accountable. The components it does require are counted in Section 2, and they are fewer than the standard formalism’s.

7 The asymmetry is dimensional

It would be easy to read the foregoing as a comparison of accuracy, and that is not the point. The two universality claims are of different *kinds*, and the dimension of the state is what makes them so.

A model whose state lives in $3N$ dimensions cannot be evaluated on a system of chemical interest, so its universality is a statement about an object that is never exhibited — an existence claim without a witness. A model whose state lives in 3 dimensions can be evaluated, so its universality is a prediction, and predictions can fail. This is the same distinction that separates a non-constructive existence proof from a computed solution, and it recurs across applied mathematics: the symbolic object is more general and cannot be inspected, the computed object is specific and can be checked. Neither is worthless, but they are not interchangeable, and universality means something different in each.

	standard QM	RealQM
state	$3N$ dimensions	3 dimensions
parameters	none, in the equation	none
cost	exponential in N	linear in N
species-blind	yes	yes
closed under composition	yes	yes
components beyond the energy	spin, antisymmetry, spin-statistics (3)	non-overlap (1)
evaluable on a protein	no	yes, 153 residues
verified range	1–3 electrons	H to Rn; H-bond geometry to 153 residues
verified precision	few parts in 10^5	few per cent; geometry better
exact in principle	yes	no
convergence ladder	yes, but unreachable	none
universality	assumed	tested

Standard quantum mechanics is verified at three electrons and asserted at 10^{23} . RealQM is verified from hydrogen to krypton on energies, and on geometry out to a 153-residue polypeptide. Over the range where either is actually used, the second is the better evidenced — a remark about evidence, not about truth.

8 What this establishes

It does not establish that standard quantum mechanics is false. Where it can be computed it is superb, and nothing here gives any reason to doubt it.

What it establishes is narrower and more useful. The conventional ranking treats the $3N$ formulation as the serious theory and a three-dimensional model as an approximation to be tolerated. On the specific question of universality that ranking inverts, because universality is a claim about scope and accountability, and only one of the two models is accountable over the scope it claims. A framework that applies in principle to everything and can be checked against nothing has a weaker title to universality than a model which applies to everything, is checked from the helium atom to a polypeptide, and agrees where it is checked. The model, its results and its code are set out at length elsewhere [27, 28, 29].

9 The tests that would decide

An argument that a model’s universality is testable owes the reader the tests. Three are set out here, each with a reference value not in dispute and each capable of going against the model. The first has already been run and the model lost it; the other two are open.

Neon’s ionization energy — run, and lost. This test is no longer outstanding. It is reported in Section 5.5: the three-dimensional angular octet on a +8 pseudo-kernel gives 89–108 eV against a measured 21.56 and a bracket of 13.33–34.48, so the attribution it was meant to confirm is refuted. The all-electron difference in three dimensions is a separate construction and does not fail this way — it gives neon $1.51\times$ the measured value, the worst of the first ten elements but the right sign and order. It is listed here because the sequence matters more than the outcome — a prediction stated, a reference not in dispute, a calculation performed, a verdict accepted. The two tests below remain open, and nothing about them is more certain than this one was.

Dispersion, by upscaling rather than by addition. Two neutral closed-shell atoms attract as R^{-6} , and the model as it stands has no such term; nor does Hartree–Fock, which gives He_2 purely repulsive, while practical density functional theory supplies the term by fitting. Adding a fitted coefficient here would place RealQM in exactly the position Section 3 criticises, and should be refused. The legitimate route is upscaling. The dispersion coefficient follows from the dynamic polarisability by the Casimir–Polder integral,

$$\alpha(i\omega) = \sum_n \frac{f_n}{\omega_n^2 + \omega^2}, \quad C_6 = \frac{3}{\pi} \int_0^\infty \alpha(i\omega)^2 d\omega,$$

and the excitation energies ω_n and transition strengths f_n are quantities this model computes. So C_6 can be assembled from the model’s own spectrum with nothing fitted and nothing added — a derived constant in the sense of Section 3, not an imported one. The construction carries its own audit: the sum must satisfy $\alpha(0) = \sum_n f_n/\omega_n^2$, which has to agree with the static polarisability computed directly. The test is then He_2 , whose well is 0.035 mHa at $R = 5.6 a_0$.

The static polarisability, which gates the other two. Polarisability is the one quantity here checked against spectroscopy rather than against another calculation, and it is also the audit on the dispersion construction. We do not have a converged value for helium: the available solvers disagree, and the number previously reported rested on a configuration carrying no electron–electron repulsion. Until that is settled nothing downstream of it can be trusted, which makes it the first of the three rather than the last.

What these have in common is worth stating, because it is the paper’s thesis in miniature. Each is a calculation the model can be made to perform, against a number it did not choose, with an outcome that can embarrass it. That is what it means for a universality claim to be testable, and it is the one thing the claim examined in Section 4 cannot offer.

References

- [1] E. Schrödinger, *Quantisierung als Eigenwertproblem*, Ann. Phys. **79** (1926) 361–376.

- [2] G. E. Uhlenbeck and S. Goudsmit, *Ersetzung der Hypothese vom unmechanischen Zwang durch eine Forderung bezüglich des inneren Verhaltens jedes einzelnen Elektrons*, Naturwissenschaften **13** (1925) 953–954.
- [3] W. Pauli, *Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren*, Z. Physik **31** (1925) 765–783.
- [4] W. Pauli, *Exclusion principle and quantum mechanics*, Nobel Lecture, Stockholm, 13 December 1946.
- [5] W. Pauli, *The connection between spin and statistics*, Phys. Rev. **58** (1940) 716–722.
- [6] P. A. M. Dirac, *Quantum mechanics of many-electron systems*, Proc. R. Soc. Lond. A **123** (1929) 714–733.
- [7] P. A. M. Dirac, *The quantum theory of the electron*, Proc. R. Soc. Lond. A **117** (1928) 610–624.
- [8] D. Giulini, *Electron spin or “classically non-describable two-valuedness”*, Stud. Hist. Philos. Mod. Phys. **39** (2008) 557; arXiv:0710.3128.
- [9] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalıy, *Classical and Quantum Computation*, Graduate Studies in Mathematics **47**, American Mathematical Society, 2002.
- [10] J. Kempe, A. Kitaev, and O. Regev, *The complexity of the local Hamiltonian problem*, SIAM J. Comput. **35** (2006) 1070–1097.
- [11] P. Hohenberg and W. Kohn, *Inhomogeneous electron gas*, Phys. Rev. **136** (1964) B864.
- [12] W. Kohn and L. J. Sham, *Self-consistent equations including exchange and correlation effects*, Phys. Rev. **140** (1965) A1133.
- [13] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*, Dover, New York, 1996.
- [14] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon, *A fifth-order perturbation comparison of electron correlation theories*, Chem. Phys. Lett. **157** (1989) 479–483.
- [15] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, *A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H–Pu*, J. Chem. Phys. **132** (2010) 154104.
- [16] C. L. Pekeris, *Ground state of two-electron atoms*, Phys. Rev. **112** (1958) 1649–1658.
- [17] S. J. Chakravorty, S. R. Gwaltney, E. R. Davidson, F. A. Parpia, and C. Froese Fischer, *Ground-state correlation energies for atomic ions with 3 to 18 electrons*, Phys. Rev. A **47** (1993) 3649.
- [18] W. Kolos and L. Wolniewicz, *Improved theoretical ground-state energy of the hydrogen molecule*, J. Chem. Phys. **49** (1968) 404–410.

- [19] J. W. Linnett, *A modification of the Lewis–Langmuir octet rule*, J. Am. Chem. Soc. **83** (1961) 2643–2653; and *The Electronic Structure of Molecules: A New Approach*, Methuen, London, 1964.
- [20] L. C. Allen, *Electronegativity is the average one-electron energy of the valence-shell electrons in ground-state free atoms*, J. Am. Chem. Soc. **111** (1989) 9003–9014.
- [21] P.-O. Löwdin, *Some comments on the periodic system of the elements*, Int. J. Quantum Chem. **3S** (1969) 331–334.
- [22] L. C. Allen and E. T. Knight, *The Löwdin challenge: origin of the $n + \ell$, n (Madelung) rule for filling the orbital configurations of the periodic table*, Int. J. Quantum Chem. **90** (2002) 80–88.
- [23] E. R. Scerri, *The Periodic Table: Its Story and Its Significance*, Oxford University Press, 2007.
- [24] D. E. Shaw *et al.*, *Anton 2: raising the bar for performance and programmability in a special-purpose molecular dynamics supercomputer*, in Proc. Int. Conf. for High Performance Computing, Networking, Storage and Analysis (SC14), IEEE, 2014.
- [25] J. Jumper *et al.*, *Highly accurate protein structure prediction with AlphaFold*, Nature **596** (2021) 583–589.
- [26] R. D. Shannon, *Revised effective ionic radii and systematic studies of interatomic distances in halides and chalcogenides*, Acta Crystallogr. A **32** (1976) 751–767.
- [27] C. Johnson, *Real Quantum Mechanics*, submitted to *Annales de la Fondation Louis de Broglie* (2026).
- [28] C. Johnson, *A 3D Multiphase Continuum Model for Atoms and the Periodic Table*, submitted to *International Journal of Quantum Chemistry* (2026).
- [29] C. Johnson, *Real Quantum Mechanics*, Body and Soul Vol. VI, in preparation; results and code at <https://github.com/claesjohnson>.