The Binding Energy Problem QCD Never Solved — and a Coulomb Model That Does

Claes Johnson[1]

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

July 20, 2026

Abstract

Quantum Chromodynamics (QCD) was put forward in 1973 as the fundamental theory of the strong interaction, and with it the accepted account of why atomic nuclei exist and hold together. Fifty-three years later, QCD has not produced a single *parameter-free, first-principles* prediction of a nuclear binding energy — not the deuteron, not the alpha particle. This is not a claim that QCD is wrong: its domain of demonstrated success is the perturbative, short-distance regime of quarks and gluons, where asymptotic freedom makes it tractable. But nuclear binding lives in the opposite, long-distance, non-perturbative regime, and there the binding energies have in practice been supplied not by QCD but by *fitted* effective theories, whose low-energy constants are read off the very nuclear data they then reproduce; lattice QCD, meanwhile, reaches light nuclei only at unphysically large quark masses, with uncontrolled extrapolation to the physical point. We document this gap, distinguish it carefully from what QCD *has* established, and then set beside it an alternative: the RealNucleus model, in which a nucleus is protons and electrons as Coulomb charge clouds, bound by the electric force alone — no strong force, no weak force. With a single scale calibrated on the deuteron and nothing else fitted, this model reproduces the alpha binding energy ($\sim$ 28 MeV, 103% of the measured value), the alpha-to-deuteron binding ratio (13.1 against a measured 12.7, a scale-free prediction), the alpha-conjugate binding ladder up to ^{40}Ca at $\sim$ 107–108%, and the emergence of saturation. The contrast is stark: the theory written to explain nuclear binding has never delivered it parameter-free, while a model containing no strong force at all delivers it with one scale. We argue this makes a long-avoided question legitimate rather than settled.

Keywords: nuclear binding energy; QCD; strong force; lattice QCD; chiral effective field theory; deuteron; alpha particle; RealNucleus; Coulomb.

1 The question

Why does a nucleus hold together? In the standard picture a deuteron is a single proton bound to a neutron, and an alpha two protons and two neutrons. The deuteron thus carries only one proton, so no Coulomb repulsion acts within it: its 2.224 MeV binding is purely the strong proton–neutron attraction. Already in the alpha, however, the two protons repel one another electrically and the nucleus holds nonetheless. Something must therefore both bind the nucleons together and, wherever protons sit side by side, overcome their Coulomb repulsion, supplying the observed binding energy: 2.224 MeV for the deuteron, 28.3 MeV for the alpha, rising to ~ 8.8 MeV per nucleon near iron.

It is worth isolating the deuteron, for it poses the question in its purest form. With only one proton there is no Coulomb repulsion to overcome; the 2.224 MeV is the strong nucleon–nucleon attraction acting alone, unobscured by electrostatics or by the many-body complications of larger nuclei. What binds a single proton to a single neutron? This is the simplest bound state in nuclear physics and the first number a fundamental theory of the strong force ought to deliver. As we record below, it has not been delivered parameter-free either — so the shortfall cannot be blamed on the difficulty of the multi-proton Coulomb problem; it is present already where no Coulomb problem exists.

The standard answer is the strong nuclear force, and its fundamental theory is Quantum Chromodynamics. The purpose of this paper is narrow and, we believe, factual: to record that after fifty-three years this theory has not predicted, from first principles and without fitting, the binding energy of any nucleus; to state precisely what that does and does not mean; and to place beside it a model that reproduces those binding energies from the electric force alone. We are not attacking QCD on its own ground. We are asking whether the ground it was supposed to cover has in fact been covered.

2 What QCD was meant to do for nuclei

QCD was formulated in 1973 with the discovery of asymptotic freedom [1, 2]: the coupling between quarks and gluons weakens at short distance and strengthens at long distance. This single structural fact organizes everything that follows. Where the coupling is weak — high momentum transfer, short distance — QCD is calculable and has been tested extensively. Where the coupling is strong — low energy, long distance, the scale of confinement and of whole nucleons bound into nuclei — it is not perturbatively calculable at all.

This is the essential point, and it is not a criticism levelled from outside: *a nucleus is a low-energy, long-distance object.* Nucleons a femtometre apart, bound by an MeV-scale energy, sit squarely in the regime where QCD's own coupling is large and its own perturbation theory fails.

The theory is at its most powerful precisely where nuclei are not, and at its weakest precisely where they are. Whatever QCD has demonstrated in the perturbative ultraviolet, the existence and binding of nuclei is a question in the non-perturbative infrared, and it is there that the account must be delivered.

Two routes have been pursued to deliver it. Neither has produced a parameter-free binding energy, and it is worth being exact about why.

3 Route one: lattice QCD

Lattice QCD evaluates the theory numerically on a discretized spacetime, and is the only route that is QCD *as such*, with no additional physical assumptions. For single hadrons — the proton mass, the low-lying meson and baryon spectrum — it has succeeded at the percent level, and this is a genuine achievement of the theory in its strong-coupling regime.

For *nuclei* the situation is entirely different. Multi-nucleon systems on the lattice are afflicted by a signal-to-noise problem that grows exponentially with the number of quarks, and by contamination from the many nearby scattering states. To make the computation tractable at all, the existing calculations of light nuclei — the deuteron, ^{3}He, ^{4}He — have been carried out at *unphysically large* quark masses, corresponding to pion masses several times the physical value (typically in the range $m_\pi \sim 300$–800 MeV against the physical $m_\pi \simeq 140$ MeV) [3, 4]. At those heavy masses a bound deuteron and alpha can be seen, but the numbers are not the physical binding energies; the extrapolation down to the physical point is long, uncontrolled, and openly disputed, and different collaborations have not converged even on whether the two-nucleon system is bound at the physical point.

The honest summary is this: fifty-three years after QCD, there exists no lattice calculation that predicts the measured binding energy of the deuteron, or of the alpha, at the physical quark masses to controlled accuracy. The one route that is QCD without further assumptions has not reached a single real nucleus.

This shows even in the pedagogy. The standard lattice-QCD textbooks close at single-hadron spectroscopy: Gattringer and Lang's *Quantum Chromodynamics on the Lattice* [6], the field's introductory text, carries the reader through gauge fields, lattice fermions, Monte Carlo, and the meson and baryon spectrum, but contains *no treatment of nuclear binding* — the deuteron and the alpha are simply absent. Multi-nucleon systems appear only in the research literature and in dedicated reviews [5], never as settled textbook material. That there is, after half a century, no textbook derivation of the binding energy of a nucleus from QCD is not an accident of exposition; it is the same gap seen from the side of what can and cannot yet be taught.

4 Route two: effective field theory — and the quiet handoff

The nuclear binding energies that theory does reproduce are obtained by a different route: chiral effective field theory (χEFT) and phenomenological nucleon–nucleon potentials [7, 8]. Here one does not solve QCD. One writes the most general interaction among nucleons and pions consistent with the symmetries of QCD, organized as an expansion in low momentum, and then *fixes the coefficients of that expansion — the low-energy constants — by fitting to nuclear and nucleon–nucleon scattering data.* With enough such constants determined from experiment, the light nuclei and their binding energies are reproduced, often well.

This is useful and it is legitimate physics. But it must be stated plainly what it is and is not. The binding energies so obtained are not *predictions from QCD*; the low-energy constants are read off the same nuclear world whose binding they then describe. The values of those constants are not, at present, computed from QCD. In effect the problem of nuclear binding has been quietly handed off from the fundamental theory to a fitted parametrization — a parametrization inspired by QCD's symmetries, but calibrated on the answer. The strong force explains the existence of nuclei in the same sense that a polynomial fitted through data points "explains" the data: it accommodates them; it does not predict them from below.

5 The gap, stated precisely

It is important to claim neither too much nor too little.

We do not claim that QCD is falsified, that it fails on quarks and gluons, or that its strong-coupling regime is understood to be empty. Its successes on the hadron spectrum are real, and belong to the regime where it is calculable.

We do claim, and take to be simply factual, that after fifty-three years:

- no calculation of QCD *as such* (lattice, at physical quark masses, with controlled extrapolation) predicts the binding energy of any nucleus, the deuteron and alpha included;
- the binding energies that theory reproduces come from effective theories whose parameters are fitted to nuclear data, not derived from QCD;
- hence the specific number that motivated the strong force in the first place — the energy that binds a nucleus against Coulomb repulsion — is still not among the numbers the fundamental theory predicts.

That a theory can be enormously successful in one regime and, for half a century, silent on the very phenomenon it was built to explain in another, is worth stating aloud rather than passing over.

6 A model that delivers the binding energies: RealNucleus

Set beside this gap a model of an entirely different character. In the RealNucleus picture [9] there is no strong force and no weak force. A nucleus is nothing but *protons and electrons as charge clouds*, bound by the ordinary Coulomb attraction — the same electric law that binds atoms and molecules, read with the charges rearranged. The neutron is a bound proton–electron pair; the deuteron is $2p + 1e$, two positive charges glued by one negative one, the nuclear cousin of the molecular ion H_2^+. The energy minimized is the plain Coulomb functional

Where the two pictures already part: the deuteron. The contrast is sharpest at the very first nucleus. In the standard picture the deuteron is a proton bound to a *neutron*, and the open question — unanswered parameter-free by QCD, as recorded above — is what binds them: what supplies the 2.224 MeV attraction between a proton and a neutral neutron. RealNucleus does not answer that question; it dissolves it. There is no elementary neutron to bind: the neutron is itself a bound $p + e$ pair, so the deuteron is not proton-plus-neutron but $2p + 1e$ — two like charges that *repel*, held together by one shared electron, exactly as one electron binds two protons in H_2^+. The binding is then no longer a mysterious strong attraction between a charge and a neutral object, but ordinary electrostatics: the electron's attraction to both protons outweighs the proton–proton repulsion. The two theories thus disagree about the deuteron already at the level of what it is made of — proton+neutron versus $2p + 1e$ — and about what holds it together — an undelivered strong force versus the Coulomb law that also binds H_2^+.

Concretely, the energy minimized is the plain Coulomb functional

$$E \;=\; \sum_i \tfrac{1}{2}\int |\nabla\psi_i|^2 \;+\; \tfrac{1}{2}\iint \frac{\rho(x)\,\rho(y)}{|x-y|}\,, \qquad \rho = \sum_i q_i\psi_i^2, \quad q_i = \pm 1,$$

over non-overlapping charge domains with free boundaries.

From this, with the *electric force only* and a *single scale* fixed on the deuteron — nothing else fitted — the model delivers:

- the **alpha binding energy**, $\sim$ 28 MeV (computed 29.1 MeV against the measured 28.3, 103%) — the very number the two QCD routes do not give parameter-free;
- the **alpha-to-deuteron binding ratio**, 13.1 against a measured 12.7 — a genuinely *scale-free* prediction, since a ratio does not see the calibration;
- the **alpha-conjugate binding ladder**, ^{4}He, ^{12}C, ^{16}O, $\ldots$, ^{40}Ca, reproduced at $\sim$ 107–108% of the measured binding, with near-constant binding per nucleon (**saturation**) emerging rather than imposed;
- the fusion reaction $d + d \to {}^4$He and alpha decay (Gamow / Geiger–Nuttall), from the same functional;
- and a result that the **electron's mass is irrelevant** to the structure: a genuinely light electron, allowed to relax, chooses a flat, charge-continuous shape, so the nuclear scale is set by the heavy proton and the atomic scale by the light electron — two sizes, one Coulomb law.

The honest caveat. This is one scale, not literally zero input: the deuteron energy sets the unit. But a unit is not a fit. Once it is chosen, every *ratio* and the *shape* of the binding-per-nucleon curve are predictions, not adjustments — in contrast to the many fitted low-energy constants of the effective-theory route. There are genuine open problems: the spin–statistics of the electron confined in the nucleus, the closed-shell structure, and the model stays deliberately silent on the neutrino. None of this is settled, and it is offered so that it can be checked and refuted.

Scope: binding, not substructure. A word on what the model does and does not address. RealNucleus resolves each nucleon and nucleus as a *structured* charge density — the solution of an electrostatic (Poisson) variational problem, with computed radii, shells, and free boundaries — and treats their *binding*. It does not address the deep-inelastic regime, in which scattering at high momentum transfer reveals the nucleon's point-like, fractionally-charged constituents (quarks and gluons); that is a separate, higher-energy question, and the correspondence between a continuum charge-density description and the parton phenomenology is left open. The claim made here concerns nuclear *binding*, not the internal parton content of the nucleon — two distinct questions at two distinct scales, which we are careful not to conflate.

The inputs: charge and size, nothing from QCD. It is worth being explicit about how little is assumed. Beyond the electric force itself and the variational principle, the model takes only two *dimensional* inputs, and both are plain experimental facts that fix the units: the unit *charge* ±1, and one *length scale* — the size of the proton. Neither is imported from QCD. The proton's femtometre size is *measured* (a charge radius, not a computed quantity), and it enters exactly as the Bohr radius a_0 enters atomic physics: as the scale that sets the units, with all dimensionless results — the alpha-to-deuteron ratio, the shape of the binding-per-nucleon curve — following as predictions. QCD would explain *why* the proton has that size; RealNucleus does not need that explanation to compute the binding, any more than atomic physics needs to derive a_0 from first principles before computing chemistry. Charge and extension are about the most primitive properties an object can have — directly measurable, theory-free numbers — and they are all the model consumes. In this precise sense RealNucleus needs neither QCD nor deep-inelastic scattering: those describe the proton's internal structure, a separate and higher-energy question that the binding problem neither requires nor contests.

7 The contrast

The two accounts can be placed side by side without rhetoric.

QCD (strong force), since 1973	**RealNucleus (Coulomb alone)**
Fundamental theory of the strong interaction; the accepted reason nuclei exist.	No strong force, no weak force; nuclei bound by the electric force only.
Calculable in the perturbative UV; nuclear binding lives in the non-perturbative IR.	The same Coulomb functional at all scales; two sizes (proton, electron) from one law.
Lattice QCD (the theory as such) reaches light nuclei only at unphysical quark masses; no controlled binding energy at the physical point.	Alpha $\sim$ 28 MeV (103%) and the ladder to ^{40}Ca at $\sim$ 107–108% from one deuteron-fixed scale.
χEFT reproduces nuclei only with low-energy constants *fitted* to nuclear data.	No fitted constants beyond the single overall scale; ratios and saturation are scale-free.

The alpha particle's $\sim$ 28 MeV is the canonical thing the strong force was invented to account for. It is reproduced, to about 103%, with a single scale, by a model that contains no strong force at all.

8 Discussion

None of this retires QCD as the theory of quarks and gluons. What it does is make a question legitimate that the field has grown used to leaving aside: if the binding energy of a nucleus can be obtained from the electric force with one calibration constant, how much of the strong-force machinery is actually required to explain why nuclei exist, and how much of what passes for explanation is fitting?

The reasonable objection — "QCD is correct, the nuclear regime is merely hard" — may well be true. But "merely hard" has meant fifty-three years without a single parameter-free nuclear binding energy, while the binding was supplied by fitted parametrizations and, separately, reproduced from scratch by a model with no strong force in it. A theory earns its standing by predicting the phenomena it was built for. On the specific phenomenon of nuclear binding, that debt is still outstanding, and the existence of a Coulomb-only account that does settle it — with its own open problems, offered for refutation — is reason to ask the question out loud rather than to assume it has already been answered.

References

[1] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Non-Abelian Gauge Theories*, Phys. Rev. Lett. **30**, 1343 (1973).

[2] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, Phys. Rev. Lett. **30**, 1346 (1973).

[3] S. R. Beane *et al.* (NPLQCD Collaboration), *Light nuclei and hypernuclei from quantum chromodynamics in the limit of SU(3) flavor symmetry*, Phys. Rev. D **87**, 034506 (2013).

[4] T. Yamazaki *et al.*, *Helium nuclei in quenched and (2+1)-flavor lattice QCD*, and related PACS-CS studies, Phys. Rev. D **86**, 074514 (2012).

[5] S. R. Beane, W. Detmold, K. Orginos, and M. J. Savage, *Nuclear physics from lattice QCD*, Prog. Part. Nucl. Phys. **66**, 1 (2011).

[6] C. Gattringer and C. B. Lang, *Quantum Chromodynamics on the Lattice: An Introductory Presentation*, Lect. Notes Phys. **788** (Springer, Berlin, 2010).

[7] S. Weinberg, *Nuclear forces from chiral Lagrangians*, Phys. Lett. B **251**, 288 (1990).

[8] E. Epelbaum, H.-W. Hammer, and Ulf-G. Meißner, *Modern theory of nuclear forces*, Rev. Mod. Phys. **81**, 1773 (2009).

[9] C. Johnson, *RealNucleus: nuclei by Coulomb forces alone*, manuscript (2026); simulations at `https://claes542.github.io/RealMolecule/gallery.html`.