RealQM for Molecules: Validation in Progress

Claes Johnson[1]

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

Version of 14 November 2026. This is a living document, updated as the validation proceeds.

Mathematical formulation by the author; implementation and numerical experiments developed with AI assistance (Claude, Anthropic).

Abstract

RealQM [1] presents the atomic world as a computable three-dimensional continuum model: electrons as non-overlapping charge densities in ordinary space, solved as the model stands. The aim is not to match the accuracy of standard quantum chemistry but to show how far such a model reaches; this is a report on work in progress. Each molecule is a partition of space into domains holding one electron or a pair, with walls placed where the density is continuous; nothing is adjusted to molecular data, and each core radius is set once per atom from its ionisation energies. Every electron pair can be carried in two parameter-free constructions, 2split and 2hom, which bracket the exact H_2 energy. Bond lengths come from the forces: twelve X–H bonds, from LiH to H_2O and the C_2 hydrocarbons, are short by 2–6%, where Hartree–Fock and DFT are within 2%. The C–C bond collapses when carbon's directions are imposed and holds, 8% short, when its walls follow the continuity rule; inner shells hold under the same rule. Polarity has the right sign wherever lone pairs are domains of their own, with sizes bracketed by the two constructions. Molecular shape is not yet produced by the model, and binding energies are too large. The failures point to one missing ingredient, the cost of pressing separate electron groups together. The model's own criteria are stated and scored, and the cost of computation, set by the grid alone, is estimated.

1 The model

RealQM [1] replaces the many-electron wave function by one-electron densities on non-overlapping domains of ordinary three-dimensional space. A molecule with nuclei $\{(\mathbf{R}_a, Z_a)\}$ is described by domains Ω_k partitioning space, each carrying $n_k \in \{1, 2\}$ electrons (two only of opposite spin) as a real amplitude u_k with $\int_{\Omega_k} u_k^2 = 1$ and density $\rho_k = n_k u_k^2$. The energy is

$$E = \sum_k n_k \Big(\tfrac{1}{2} \int_{\Omega_k} |\nabla u_k|^2 - \int_{\Omega_k} K u_k^2 \Big) + \sum_k \tfrac{n_k(n_k-1)}{2} U_{kk} + \sum_{k<l} n_k n_l U_{kl} + V_{KK}, \tag{1}$$

with $K = \sum_a Z_a/|\mathbf{x} - \mathbf{R}_a|$, U_{kl} the Coulomb interaction of the unit densities u_k^2 and u_l^2, and V_{KK} the nuclear repulsion. A pair domain ($n_k = 2$) counts its single pair once.

What to expect of a basic model. For one electron the energy (1) is Schrödinger's equation for the hydrogen atom, and it is exact there. That exactness is not a test: the equation was formed to reproduce the hydrogen spectrum, and its success there is a calibration. It can give the impression that the model is not only basic but nearly perfect. Being basic carries no such promise beyond the case it was fitted to. From H_2 on, the model is an extension of a one-electron equation to many electrons by one assumption, non-overlapping domains with free walls, and nothing in that assumption says the result should be as accurate as for hydrogen. The deviations reported below are therefore not surprising; they measure what the basic form leaves out. The standard extension of Schrödinger's equation to $3N$ dimensions is not tested by the hydrogen spectrum either; it is tested separately, for a few electrons, where it agrees with observation, and it is viewed as exact beyond them (see Testability below).

In accuracy RealQM does not compete with the computable methods of standard quantum chemistry. Standard quantum mechanics is a calculus of observable outcomes; whether it is also a model of physical reality is an open question, since its candidate reality, a wave function in $3N$-dimensional configuration space, has no agreed physical interpretation. It predicts physics without settling what the physics is. RealQM takes up that open question: it seeks a model in which the physics happens in three-dimensional space, and it is to be judged by how much of the established predictions such a model reproduces. Its case therefore rests on a physical picture in three dimensions, nothing fitted, and a cost set by the grid alone (Section 8).

What is observable. The two formulations predict the same kinds of measured quantity: energies, geometries, forces, dipoles, and the total electron density $\rho(\mathbf{x})$ that X-ray and electron diffraction measure. Standard quantum mechanics obtains ρ by integrating its $3N$-dimensional wave function over all electrons but one; RealQM obtains it by adding its domain densities. RealQM is the more directly interpretable of the two, since every object it uses is a density in ordinary space, of the same kind as what diffraction measures, while the wave function in configuration space is not. Its individual domains are not themselves measured; diffraction sees only their sum. The total density offers a test independent of energies. Accurate diffraction gives deformation densities, the change of ρ on bonding, for small molecules and crystals; RealQM predicts how much charge sits in the bond region, and the over-binding found here suggests too much.

Testability. Standard quantum mechanics is often presented as a perfect predictor within its realm. That is established where it can be computed exactly; for many electrons it is an extrapolation that cannot at present be tested. It is unrefuted wherever it can be tested exactly, a few electrons with relativistic and radiative corrections, and it is the reference for accuracy. For many electrons its equation cannot be solved, and a disagreement between a computed result and an observation is assigned to something other than the equation: first to the approximation (basis, functional, level of correlation, convergence), then to physics left out of the calculation (relativity and spin–orbit coupling for heavy atoms, vibrational averaging and temperature, the environment), then to the measurement or its interpretation. The assignment is often justified, since results move toward observation as approximations are refined, but it means that no result for many electrons can count against the equation itself; its correctness there is an extrapolation. RealQM is computed as it stands, with nothing between model and result but the mesh, so each deviation reported here counts against the model. The comparison is therefore not symmetric: one model is protected by being uncomputable, the other exposed by being computable. RealQM is offered as a model with stated, measurable errors, computed as it stands: less accurate, but open to test in every case.

The question behind the model. How can a linear equation in $3N$-dimensional configuration space, a formal object with no physical meaning in the three-dimensional space where atoms are, predict the physical atomic world as precisely as it does where it can be tested? The question needs an answer. "Shut up and calculate" declines it, and the proposed answers move it rather than settle it: to take configuration space itself as physical gives up the three-dimensional world as the arena of physics; to read the wave function as a calculus of probabilities leaves unexplained why that calculus works; a reformulation in terms of the three-dimensional density (Hohenberg–Kohn) adds no information, since its functional is defined through the $3N$-dimensional problem. RealQM takes the question as open and answers it constructively: it seeks a model in three-dimensional space, with a physical picture, that is computable as it stands, and measures how much of the established predictions such a model reproduces. The results here are that measure for the present model. Without an answer, the success of standard quantum mechanics is predictive, not explanatory: it says what will be observed, not what the atomic world is. A century of effort has produced no agreed account of what the formalism describes or why it works, only many interpretations, sharpened by Bell's theorem and its experimental tests and by decoherence, none accepted. Together with the incompatibility of quantum theory and gravity, this is regarded as one of the central open problems of modern physics, and RealQM approaches it by seeking a computable model with a physical picture in three dimensions.

Boundary conditions. Each u_k satisfies a homogeneous Neumann condition on $\partial\Omega_k$: the kinetic energy is counted only inside the domain, and the amplitude at the wall is free. The boundaries are placed by *density continuity*, $\rho_k = \rho_l$ across each interface. Energy minimisation over the boundary position is not used: in this model a ragged or displaced boundary can lower the energy (Section 5.3), and only continuity holds it.

2 Atoms

Atoms are computed in spherical symmetry, with all electrons organised in non-overlapping radial shells following the Atom Simulator configurations [2] (Li 2+1, B 2+3, ...): a one-dimensional problem in the radius, with each shell boundary free and placed by the same continuity rule as in three dimensions. For use in molecules the atom is then reduced to a *pseudo-atom*: the $1s^2$ core (and further inner shells) is frozen into a pseudo-nucleus of charge Z_x with a hard core of radius r_c, and only the outer shell is carried in three dimensions, split into domains of at most two electrons (boron: pseudo-B +3 with an outer 2hom pair and one electron). A radial shell holding three or four electrons is a shared density in one dimension, which no domain may be in three dimensions; it stands for the spherical average of its angular split into domains of at most two (boron's 3 as 2_{hom}+1, a 4 as a tetrahedron of single electrons). Shells hold at most four—the pair and the tetrahedron, an octet being 4+4, never a single shell of six or eight. Energy does not select the configuration: more sharing always lowers the one-dimensional energy (Li as one shell of 3 is 13.7% too deep, as 2+1 within 0.5%), so the configuration is set by the Pauli cap of two per domain, not by minimisation. The core radius is a property of the atom, not of any molecule, and it is obtained in two ways. *Fitted*: the radius at which the pseudo-atom alone, in three dimensions, reproduces the sum of the ionisation energies of the electrons it carries. *Computed*: the core–valence interface of the one-dimensional all-electron atom, which for a $1s^2$ core is fixed by the model itself, independent of the box:

	r_c fitted to IE (3D)	r_c computed (1D, boxes 6–20 a_0)	difference
Li (2+1)	1.99	1.98–2.05	within 3%
B (2+3)	1.08	0.92–1.00	8–17%
Be (2+2)	1.37	pending	
C (2+4)	1.0	pending	

For Li the two agree, and the core radius is an output of the model. For B the computed radius is shorter than the fitted one. We use the fitted radius throughout, and state the computed one beside it; an outer core shell beyond $1s^2$ (Na and below) has no box-independent interface in one dimension, and there only the fitted radius exists. At $h = 0.1$ the hard core excludes whole cells, so the fitted radius is defined to a few hundredths of a_0.

3 Molecules by successive reduction

Molecules are built by reducing the full model step by step, each step tested on systems where the step before it can still be run.

1. *All electrons* (H_2, He): every electron in its own domain or in a pair, with no pseudo-nucleus. Here an electron pair has two constructions to choose from, 2split and 2hom (Section 5).

2. *Frozen* $1s^2$ *core*: for every molecule beyond H_2 the atoms enter as pseudo-atoms (Section 2); only valence electrons are computed, in domains of at most two, with the inner boundary r_c fixed and all outer boundaries free. A free boundary charges nothing for crowding, so the core sphere must be fixed (Section 6.1). LiH and BH are computed this way.

3. *Reduced dimension* for linear molecules whose boundaries are fixed planes—the midplane of a homonuclear X_2, by symmetry, or planes through the nuclei, by construction: the axisymmetric solver (Section 4).

4. *Large molecules*: pseudo-atoms with reduced valence charges on a coarse mesh. The same `molecule.js` has been run on a 238-residue protein (green fluorescent protein, 2307 atoms, 300^3 grid, $h = 0.4\,a_0$) and on crambin (540 atoms, 200^3, $h = 0.25\,a_0$) on the GPU of a laptop in a web browser. This shows that large molecules are within computational reach; it is not a validation. At this mesh the bond is not resolved, the folds shown use contact restraints, and at this size the solver uses the total-Poisson path for the electron repulsion (Section 4), which must be repaired before results there can be judged.

4 Solvers, and the check between them

`molecule.js`. One amplitude field on a Cartesian grid, each cell labelled by its owning domain, so the domains are disjoint by construction. The kinetic stencil omits every face joining two labels (true Neumann). Boundaries move by label flips: a boundary cell changes owner when the neighbouring domain's density exceeds its own. The Coulomb potential of the other electrons is a direct per-electron Poisson solve. The nuclear kernel is clamped, $-Z/\max(r, R_{\text{sing}})$, with $R_{\text{sing}} = 2h$ unless pinned. The amplitude is relaxed in imaginary time; convergence is declared when the mean energy over a window of 3000 steps changes by less than $10^{-4}\,E_h$ between successive windows, since a free boundary keeps flipping single cells indefinitely and a step-to-step test either never fires or fires early.

Axisymmetric reduction. For a homonuclear X_2 the domain boundary is the midplane by symmetry, and the model reduces exactly to (z, ρ) with a plane boundary. The same holds for any linear molecule whose walls are all fixed planes, such as planes through the nuclei separating the halves of a pseudo-atom's valence, with every bond a 2hom pair (Section 6.3). We use it only there: a free wall between unlike atoms curves, and a plane would impose it. The Neumann condition is imposed by assembling the stiffness only over faces inside a domain. Each domain's lowest state is an exact sparse eigensolve, and the potentials are iterated to self-consistency. Mesh ladders down to $h = 0.05$ are cheap.

The check. The electron–electron energy of the two solvers agrees to 1.4%, and the total energies to the extent their kernels allow. At $R = 2.0\,a_0$, $h = 0.1$ and box $10\,a_0$:

	$E(H_2)$	$2E(H) - E(H_2)$
`molecule.js` (3D, free boundary)	−1.20734	0.2251
axisymmetric (2D, exact eigensolve)	−1.23685	0.2241
exact	−1.1381	0.1381

The totals differ by the kernel (clamped at $2h$ in 3D, analytic in 2D), which largely cancels in the binding energy; the binding energies agree to 0.5%. The three-dimensional solver reproduces the model.

An independent energy check. The agreement above is for two domains. With more than five, `molecule.js` by default replaces the per-electron Poisson solves by one total Poisson solve with a self-interaction correction, and that path had not been checked. We therefore recompute every energy from the converged state alone: the amplitude and the domain label of each cell are exported, and V_{ee} is recomputed by a zero-padded FFT convolution with $1/r$ (each pair of domains once, and $(n-1)/n$ of a pair domain's own Hartree energy), V_{eK} from an analytic kernel, T from finite differences inside each domain with free walls, and V_{KK} as a point sum. On the H atom this agrees with the solver to 10^{-5}, on H_2 to $0.008\,E_h$. On CH_4 it exposed the total-Poisson path: for its own converged density the true V_{ee} is 10.28 where the path reports 7.81, so it reported an energy $2.5\,E_h$ too low and also converged to a worse density (true energy −9.00 against −9.27 on the per-electron path). All results here use the per-electron path. That path over-counts V_{ee} slightly, by $0.04\,E_h$ for the C atom to 0.36 for ethane, growing with the number of pairs; every energy quoted for more than two domains is the independently recomputed one.

5 Hydrogen and helium: two constructions of a pair

5.1 Two constructions of an electron pair

RealQM carries no spin variable; spin enters only through the Pauli principle, as the rule that a domain holds at most two electrons, and two only if of opposite spin. Electrons in different domains are excluded from each other by disjointness, whatever their spin. A pair of opposite spin can therefore be carried in two parameter-free ways. In the *2split* construction each electron owns its own domain and the two meet at a free wall: their avoidance is complete and permanent. In the *2hom* construction both electrons share one domain and one density, the overlay that Pauli permits for opposite spins, each feeling half its Hartree field: there is no avoidance beyond the mean field, and the pair is a restricted Hartree–Fock pair. A one-dimensional radial check on Li^+ makes the last point exact: the 2hom $1s^2$ pair gives −7.23630 against the Hartree–Fock −7.2364. The true state avoids partially and dynamically.

5.2 H_2 and He

For both two-electron systems the exact value lies between the two constructions (H_2 axisymmetric, $R = 1.4$, $h = 0.1$):

H_2	T	V_{eK}	V_{ee}	D_e
2split, two one-electron domains	1.291	−3.843	0.545	0.281
exact	1.1745	−3.650	0.587	0.1745
2hom, one pair domain	1.153	−3.679	0.664	0.138

He	E
2split (two hemispheres)[a]	−3.026
exact	−2.9037
2hom (one pair, 1D radial)[b]	−2.845

[a] Axisymmetric; the pair repulsion is 4% below exact. [b] The Atoms paper [2]. (For H_2 the exact T is $-E$ by the virial theorem, and V_{eK} follows.) In both cases the exact value lies strictly between the two constructions, with no adjustable quantity on either side: $D_{2\mathrm{hom}} < D_e < D_{2\mathrm{split}}$ and $E_{2\mathrm{split}} < E < E_{2\mathrm{hom}}$. This is the model's energy statement. An intermediate construction that reproduced the exact value would be a fit.

The two constructions are distinct cases with no bridge between them. A pair either has a wall through it or it does not, and no parameter of the model carries one into the other. Charging the wall (Section 5.3) does not do it: as β grows the density moves away from the wall, not towards the shared pair. The situation is like a bifurcation. The pair sits on one branch or the other, and the states in between are not solutions of the model. The exact value lies between the two branches in energy, but it is not on a path that joins them.

The geometry is sharper than the energy. The 2split H_2 converges in the axisymmetric reduction to $R_e = 1.368\,a_0$, short of the exact 1.401 by 2.3% and flat to 0.2% under a 2.5-fold mesh refinement.

What makes the H_2 bond is not agreed among experts, even though the exact energy is not in doubt. One account, from Slater [9] and Feynman [10], is electrostatic: charge gathers between the nuclei, and the Hellmann–Feynman force on each nucleus pulls it inward. The other, from Hellmann [11] and Ruedenberg [12], puts the cause in the kinetic energy: an electron spread over both atoms lowers its kinetic energy, and the potential energy follows. The two accounts are still argued today [13]. Both read the same virial balance, $\Delta T = -\Delta E$ and $\Delta V = 2\Delta E$ at equilibrium, which holds whichever is the cause. The model takes no side by assumption. Its 2split bond is made by charge held in the bond region at no kinetic cost (Section 5.3), so it is closer to the electrostatic account. That is a property of the construction, not a settlement of the question.

That the cause of the simplest chemical bond is still argued, about ninety years after its energy was first computed, is not a failure of computation. The exact energy is not a fit: no coefficient is adjusted to experiment, and it corrected the measured dissociation energy rather than following it [16, 17]. Standard quantum mechanics explains the bond by law: the energy follows from the Schrödinger equation. It does not explain it by mechanism. The solution is a function on six-dimensional configuration space; a division of its energy into kinetic and potential causes depends on the reference state and the path chosen; and what happens in three dimensions is read into it, differently by different readers. Schrödinger asked for the second kind of explanation. A model in three dimensions carries a mechanism directly: the 2split bond is charge held in the bond region at a free wall that costs no kinetic energy. Its energy is not exact either, so the model does not settle the dispute; it shows what an explanation that can be checked looks like.

5.3 What places the 2split side

The 2split energy is not an error in the solution: two independent solvers agree, the energy converges, and every term is counted. Its components show what the free wall does. Against exact, the nuclear attraction is larger (centroid of electron 1 at -0.785, the nucleus at -0.70) and V_{ee} smaller: a Neumann wall lets each electron keep full amplitude right up to the midplane—in the bond region, where both nuclei attract it—at no kinetic cost, and the other electron, confined strictly to the other side, never competes for that region. Charging the wall shows this directly. With a surface energy $\frac{\beta}{2}\int_{\partial\Omega_k} u_k^2$ on the interface:

β	T	V_{eK}	V_{ee}	centroid
0	1.291	-3.843	0.545	-0.785
0.15	1.277	-3.803	0.526	-0.832
0.5	1.265	-3.695	0.485	-0.936
2	1.281	-3.255	0.381	-1.269
10	1.238	-2.607	0.288	-1.713

As soon as the wall costs anything the density leaves the interface for the far side of its own nucleus and the nuclear attraction falls, with the kinetic energy almost unchanged. The free cut at the wall is the model's central assumption; it sets the 2split energy, while the 2hom pair has no wall.

5.4 Two states of a pair

The two constructions are two different states of the pair, and standard quantum mechanics has both. The 2hom pair, two electrons of opposite spin carrying one density over one domain, is the restricted Hartree–Fock singlet: one doubly occupied orbital. In H_2 at $R = 1.4$ the axisymmetric 2hom energy is -1.133585 at $h = 0.1$, against the Hartree–Fock limit -1.133630 [15]. The two share the known Hartree–Fock behaviour: H^- is unbound, and the pair does not dissociate into two neutral atoms, because half of it stays ionic. The 2split pair is the Heitler–London state. Each electron keeps to its own region, which is avoidance in the extreme, and spin plays no part: without an exchange term the singlet and the triplet are the same. It dissociates correctly and binds H^-. The exact ground state is neither. It is a correlated singlet, which standard quantum mechanics builds by mixing the shared and the separated configurations. A real density with one construction per pair has no such mixture, and that is why the exact energy lies between the two and is reached by neither (Section 5).

In Hartree–Fock the two states are joined by a bifurcation. The restricted solution is stable near equilibrium and becomes unstable beyond the Coulson–Fischer point, about $R = 2.3$, where it breaks symmetry into the unrestricted solution, which separates the electrons like 2split [18]. The model has the two branches but not the path between them. Domains are disjoint, and two electrons share a domain only with the same density, so the partial overlap through which the Hartree–Fock solution breaks symmetry is not available. A 2hom pair therefore cannot decay into a 2split pair. It is also lower in energy at no R:

R	D(2hom)	D(2split)	$E_{\mathrm{split}} - E_{\mathrm{hom}}$
1.0	0.0847	0.2305	−0.1458
1.2	0.1246	0.2679	−0.1434
1.4	0.1332	0.2747	−0.1416
1.6	0.1259	0.2663	−0.1404
2.0	0.0912	0.2308	−0.1396
2.4	0.0489	0.1894	−0.1406
3.0	−0.0111	0.1336	−0.1447
4.0	−0.0887	0.0679	−0.1566
5.0	−0.1410	0.0308	−0.1718
6.0	−0.1755	0.0114	−0.1869
8.0	−0.2135	−0.0037	−0.2098

(Axisymmetric, $h = 0.1$, $D = 2E_{\mathrm{H}} - E$ on the same mesh; 2hom mirror-symmetric.) The two curves do not cross: 2split is lower at every R, by 0.14 near equilibrium and by more as the bond stretches. If a path existed, energy minimisation would take every pair to 2split, including a bonding pair at equilibrium, where 2hom is the closer one to experiment (Section 5.5). The free wall over-favours avoidance. So the construction of a pair is not chosen by energy. It is chosen by the physical rule of the next subsection: shared pairs are 2hom, and electrons that must avoid each other are 2split.

Breaking a bond requires 2split. The scan shows the division of labour directly. The 2hom pair stops binding near $R = 3$ and rises steeply beyond it, 0.21 above two atoms at $R = 8$, because half of the shared pair stays ionic, as in restricted Hartree–Fock. Left without the mirror symmetry, a stretched 2hom pair does not turn into a 2split pair. It collapses onto one nucleus, the ionic state $\mathrm{H}^- + \mathrm{H}^+$ (97.5% of the density on one side at $R = 6$), which is a higher solution than the symmetric pair. Only 2split goes to two neutral atoms. Its tail falls off too slowly, however: it still binds 0.011 at $R = 6$, where the exact binding is nearly zero, and it is slightly repulsive (0.004) at $R = 8$. This matters for the later aim of reactions between fragments. A reaction breaks or forms bonds, so somewhere along its path a pair has to change from 2hom to 2split, and the model has no path between them. The energies at the two ends of a reaction can be computed in the model; a barrier between them needs a rule for where the pair changes, which the model does not yet have.

The natural rule from inside the model is to switch where 2hom stops binding, $R \approx 2.95$. The two curves do not meet there (Figure 1): at the switch 2hom has $D = 0$ and 2split $D \approx 0.14$, so the energy drops by $0.14\,E_h$ on stepping outward, more than the whole exact binding at that distance (0.06), and the combined curve has its lowest point at the switch, a well that would hold a stretched molecule instead of letting it dissociate. The rule classifies a pair as bonded or broken and describes the two end states, but no curve through the switch. For comparison the figure includes the pair of standard quantum mechanics between the two constructions, computed on the same grid: two overlapping one-electron functions coupled as a singlet with the exchange term, $E = (Q + A)/(1 + S^2)$ (generalised valence bond). The energy alone chooses its overlap, $S = 0.80$ at $R = 1.4$ falling to 0.06 at $R = 6$, and the curve follows the exact one to within $0.02\,E_h$ at every R ($R_e = 1.446$). Exchange needs indistinguishable electrons, which the model does not have; within it, the two constructions bracket the curve at every R and nothing selects a state between them.

A bridge by competition. Reaction–diffusion systems suggest a bridge. Two densities that diffuse and compete through a term $\kappa u_1^2 u_2^2$ separate into disjoint regions as $\kappa \to \infty$, the strong-competition limit [19]. We give each electron its own density over the whole space and add $\kappa \int u_1^2 u_2^2$ to the energy. At $\kappa = 0$ this is unrestricted Hartree–Fock for two opposite spins: the Hartree

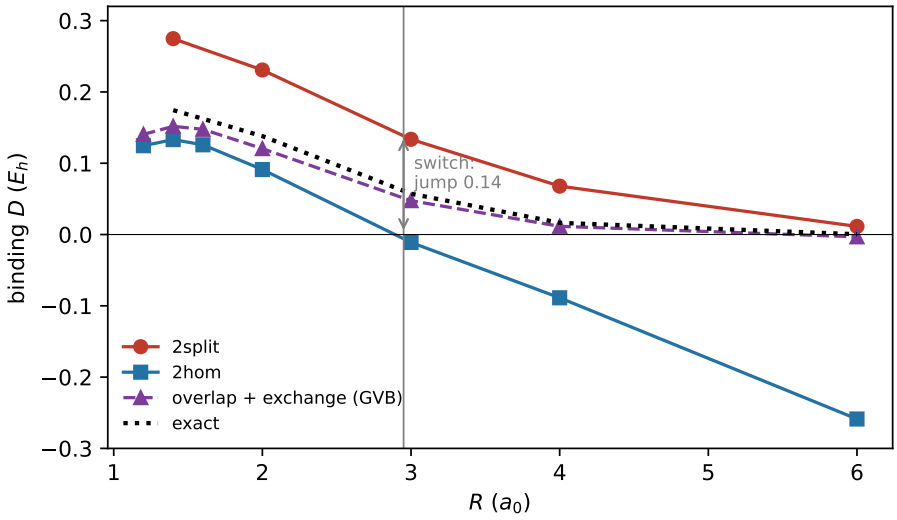


Figure 1: H_2 binding $D = 2E(\mathrm{H}) - E(\mathrm{H_2})$ against R (axisymmetric, $h = 0.1$): 2split, 2hom, the overlapping singlet pair with exchange (generalised valence bond), and the exact curve (literature values [14], quoted, not recomputed). The grey line marks a switch from 2hom to 2split where 2hom stops binding; the arrow is the energy jump there.

repulsion between two orbitals, with no exchange term. Started each time from the electrons on opposite sides (axisymmetric, $h = 0.1$):

R	κ	D	orbital overlap	state
1.4	0	0.1332	1.000	symmetric = 2hom
1.4	5	−0.0429	1.000	symmetric
1.4	6	−0.0707	0.973	broken
1.4	1000	−0.5604	0.045	segregated
2.2	0	0.0702	1.000	symmetric
2.4	0	0.0497	0.923	broken
4.0	0	0.0012	0.265	broken

The family has a pitchfork bifurcation. At $\kappa = 0$ the symmetric solution breaks between $R = 2.2$ and 2.4, which is the Coulson–Fischer point. At $R = 1.4$ it breaks between $\kappa = 5$ and 6. Near both breaks the iteration slows down, as it does at a bifurcation. In the (R, κ) plane a curve of symmetry breaking separates the shared pair from separated electrons, and the broken $\kappa = 0$ branch dissociates correctly ($D = 0.001$ at $R = 4$). The far end of the family, however, is not 2split. Segregation by competition forces the densities to vanish at the interface, and this hard wall costs energy: at $R = 1.4$, $\kappa = 1000$ is 0.56 above two atoms, while 2split binds by 0.275. No overlap penalty reaches the free Neumann wall. The wall is the model's own assumption and the source of its over-binding (Section 5.3), and 2split stands outside the family of overlapping states. A switching rule by competition is therefore available, but it switches into the hard-wall branch. Using it for reactions would add these states to the model rather than join its two constructions.

5.5 Beyond two electrons

The same two constructions apply to any pair in a molecule. For LiH, with Li as a +1 pseudo-nucleus ($r_c = 1.99$ from its ionisation energy) and the bonding pair as 2split (Li valence and H

electron in separate domains) or 2hom (one pair over Li^+ and a bare proton, Section 6.1), D_e is 0.063–0.073 (2split) and 0.095 (2hom) against the exact 0.092. For BH (pseudo-B +3 with a 2hom lone pair) D_e is 0.276 (2split) and 0.155 (2hom) against 0.135 (Table 1). In both, the 2hom bonding pair is the closer.

The picture also tells where each construction belongs. A pair that shares one region—a core $1s^2$, or a bonding pair near equilibrium—is naturally 2hom. Electrons that must avoid each other need 2split: the 2hom H^- is unbound, exactly as in Hartree–Fock, while the 2split H^- binds; and a 2hom pair cannot dissociate a bond into two neutral atoms. We report both, and choose by this rule.

The results below divide the roles as follows. *Geometry is common to both cases:* where both were run, the bond lengths differ by a few per cent at most (Li_2 5.1/5.15, N_2 1.95/1.95, LiH 2.85/2.98, BH 2.29/2.38 (solver energy, $h = 0.1$); Table 1 and Section 6.3), so the bond length does not depend on the choice. *2split is the general case and over-binds:* its D_e is above experiment in most molecules computed, up to about twice (BH 0.276 against 0.135, OH 0.371 against 0.170). The exceptions are LiH (0.063–0.073 against 0.092) and B_2 (at experiment). *2hom is the shared pair:* two opposite spins in one domain. It is the closer energy for a shared pair at equilibrium (LiH 0.095 against 0.092, C_2 0.242 against 0.232, N_2 0.370 against 0.364) and under-binds where the electrons should avoid each other. It also gives the LiH dipole, of the right sign and 11% high, where one electron per domain inverts it. Dipoles in general are not settled by either case. In BH only a cone-shaped boundary corrects the dipole, at the cost of bond length and binding (Section 6.2). In OH both cases are about twice too large. H_2O needs split O for the right sign.

6 Results

The molecules below are a first selection of basic examples, chosen where the model's constructions are settled; many more are to follow, and those tried beyond the selection are collected in Section 7. The criterion is geometry: bond lengths. Binding energies are listed as a diagnostic of the construction, not as a target. **[In progress.]** Bond lengths from an R-scan at three meshes ($h = 0.10, 0.07, 0.05$) per bond with `molecule.js`, extrapolated in h, with the kernel clamp pinned so that the ladder refines a fixed problem; axisymmetric values beside them for X_2 only. Energies for both constructions, 2hom and 2split, at the computed geometry.

		bond length (a_0)				binding D_e (E_h), diagnostic		
molecule	construction	3D	axisym	exp.	error	2hom	2split	exp.
H_2	1+1	pending	1.368	1.401	−2.3%	0.138	0.281	0.1745
LiH	Li^+ (r_c) + pair	2.85	—	3.015	−5%	0.095	0.063–0.073	0.092
BH	B 2+3 (2_{hom}+1)	2.23	—	2.329	−4%	0.155	0.276	0.135
OH	O +3 (2_{hom}+1)	1.79	—	1.832	−2%	0.241	0.371	0.170
BeH_2	Be $2+2_{split}$	2.44	—	2.52[a]	−3%	0.248[b]	n.c.	0.2231

Table 1: Bond lengths and the 2hom and 2split binding energies against experiment. H_2 energies are axisymmetric ($R = 1.4$, $h = 0.1$); LiH is three-dimensional (box $20\,a_0$, $h = 0.1$), its bond length from the 2hom construction (2split: 2.98–2.99). BH is three-dimensional (box $12\,a_0$, pseudo-B $r_c = 1.08$), $R = 2.1, 2.3, 2.5$; bond length from 2hom, fitted to the independently recomputed energy (Section 4) on a mesh ladder: 2.273, 2.248, 2.232 at $h = 0.1, 0.07, 0.05$, still shortening (the earlier 2.29 was a fit to the solver's own energy at $h = 0.1$; 2split, solver energy: 2.38). BeH_2: pseudo-Be $r_c = 1.37$, box $12\,a_0$, $R = 2.3, 2.5, 2.7$; [a]gas-phase r_0 (vibrationally averaged); [b]atomisation energy (two bonds), experiment $0.2231\,E_h$ [6]; n.c.: the 2split runs do not converge (the two sides trade cells and the symmetry breaks). OH: pseudo-O +3 (only the outer 3-shell carried, as 2hom lone pair +1 behind a plane), $r_c = 0.59$, $R = 1.65, 1.83, 2.0$; 2split $R_e \approx 1.69$; dipole 3.6 D (2hom), 2.7 D (2split) against 1.66. The axisymmetric reduction is used for X_2 and for linear molecules with fixed walls only.

6.1 LiH: pseudo-Li and one bonding pair

LiH is the simplest polar bond with an inner shell, and is viewed as ionic (Li^+H^-) rather than covalent: its dipole, 5.88 D, is 77% of that of point charges ±1 at $3.015\,a_0$ (7.66 D). The model reduction reflects this: the Li $1s^2$ core is excluded by a fixed sphere of radius $r_c = 1.99\,a_0$ around the Li nucleus, which then carries charge +1 (pseudo-Li); the two valence electrons form one 2hom pair over pseudo-Li and a bare proton. The radius is the one at which the pseudo-Li atom alone reproduces the Li ionisation energy; the one-dimensional all-electron Li atom places its own core–valence interface at 1.98–$2.05\,a_0$ over boxes 6–$20\,a_0$, within 3% of it. All runs are `molecule.js`, $h = 0.1$, box $20\,a_0$; binding energies are $E(\mathrm{Li}) + E(\mathrm{H}) - E(\mathrm{LiH})$ with both atoms built the same way and in the same box.

R	E	$D(R)$	μ (D)
2.6	−0.76776	0.0880	—
2.8	−0.77436	0.0946	6.45
3.0	−0.77259	0.0928	6.53
3.2	−0.76709	0.0873	6.66
3.4	−0.75999	0.0802	6.82
exp.		0.092 at 3.015	5.88

A fit through the five points gives $R_e = 2.85\,a_0$ (2.81–2.89 over three-, four- and five-point fits), 5% short of experiment, and $D_e = 0.0947\,E_h$, 3% over. At $R = 3.0$ the dipole is 6.53 D, of the right sign (Li^+H^-) and 11% high; it is 86% of the point-charge value at that R, so the pair, free to distribute itself over both centres, settles largely on H. Nothing was adjusted to LiH. The 2split on the same pseudo-Li (Li valence and H electron in separate domains) gives a slightly longer bond, $R_e = 2.98$–2.99 (−1%), and under-binds, $D_e = 0.063$–0.073 (Section 5.5). The core is a hard exclusion with a radius set by the atom: a frozen-core model, stated as such. A 2hom core held

inside the same fixed sphere is equivalent for the valence, since a spherical charge inside the sphere acts outside it as a point charge. The sphere must be fixed: a free boundary between core and valence charges nothing for crowding, and with it the valence presses into the core and $E(R)$ falls monotonically as H approaches.

6.2 BH and BeH_2

BH. Pseudo-B +3 carries its outer 2+3 shell as a 2hom lone pair and one electron, the two separated by a fixed surface through the B nucleus; the B|H boundary is free. Two divisions of the 3 were run, each with its own r_c fitted to the B ionisation energies (box $12\,a_0$, $h = 0.1$): a plane perpendicular to the bond (each domain half the solid angle, $r_c = 1.08$), and a cone about the bond holding one third of the solid angle for the single electron and two thirds for the pair (half-angle 70.5°, $r_c = 1.11$), in proportion to their electrons. For the bond, 2hom (one pair over the B front and a bare proton):

2hom BH	R_e	D_e	μ (D)
plane, $\frac{1}{2}$:$\frac{1}{2}$	2.29	0.155	3.5
cone, $\frac{1}{3}$:$\frac{2}{3}$	2.17	0.224	1.09
exp.	2.329	0.135	1.27

The proportional division corrects the dipole, by moving charge behind B, and costs bond length and binding: the pair confined to the cone draws the proton in. The plane gives bond and energy, the cone the dipole; neither gives all three. (These R_e are fits to the solver's energy at $h = 0.1$, which puts both constructions on the same footing; for the plane the independently recomputed energy gives 2.273 at $h = 0.1$, shortening to 2.232 at $h = 0.05$, Table 1.) 2split binds more strongly in both (0.276 and about 0.34), with dipoles 4.0 and 1.2–1.3 D. Table 1 uses the plane.

BeH_2. Pseudo-Be +2 (2+2, $r_c = 1.37$) carries its two valence electrons as single electrons on either side of a fixed plane through Be, each facing one H. With each bond a 2hom pair over a Be half and a bare proton, $R_e = 2.44$ against the gas-phase $r_0 = 2.52$ (−3%), and the atomisation energy is $0.248\,E_h$ against 0.2231 [6] (+11%). The 2split construction does not settle: the two H domains and the two Be halves trade cells without end, the symmetry breaks (a residual dipole of 0.2 D), and the convergence window never closes.

6.3 Homonuclear X_2

For X_2 the axisymmetric solver is used. Each X is a pseudo-atom whose inner shells, as given by the Atom Simulator configuration, are frozen into a core of charge Z_x with r_c fitted to the sum of the ionisation energies of the electrons carried; the outer shell lies along the axis as a domain behind X (a 2hom lone pair, or one electron) and one bonding electron, separated by a fixed plane through X. The two bonding electrons form either a 2split pair (the midplane) or one 2hom domain between the nuclei. Only a σ pair is representable. R-scan 0.8–$1.3\,R_{\mathrm{exp}}$, $h \approx 0.1$ (adjusted per R to keep the nuclei on cell faces, a jitter of about $\pm 0.005\,E_h$):

X_2	configuration	r_c	R_e: 2hom	/ 2split	D_e: 2hom	/ 2split	exp. R_e / D_e
Li_2	2+1, +1	2.00	5.1	5.15	0.021	0.054	5.05 / 0.038
Na_2	2+4+4+1, +1	2.16	5.5	5.55	0.020	0.050	5.82 / 0.028
B_2	2+2+1, +1	0.88	2.87	2.9	0.045	0.109	3.005 / 0.110
C_2	2+2+2, +2	0.91	2.55	2.58	0.242	0.322	2.348 / 0.232
N_2	2+3+2, +2	0.59	1.95	1.95	0.370	0.480	2.074 / 0.364
O_2	2+3+3, +3	0.59	2.0	2.05	0.349	0.48	2.282 / 0.192
F_2	2+4+3, +3	0.49	no minimum		> 0.35	/ > 0.52	2.668 / 0.061

For Li_2, Na_2, B_2, C_2 and N_2, cores +1 and +2, the experimental binding lies between the two constructions and the bond length is within 10%; B_2 and N_2 require the Atom Simulator configurations (as pseudo-+3 with a lone pair, B_2 over-binds twofold). O_2 and F_2 are not reproduced in any configuration tried (Section 7).

Transfer of the carbon model: C_2H_2. The C_2 carbon (pseudo-C +2, $r_c = 0.91$, valence 1+1) carried unchanged into acetylene, with every bond a 2hom pair (each C–H pair the H electron and the outer C half over a bare proton; the C–C pair the two facing halves), has all its walls fixed at the planes through the C nuclei, so the axisymmetric solution is that of the three-dimensional construction. Each bond scanned with the other at its observed length:

C_2H_2	C–C	C–H	atomisation at obs. geometry
2hom	2.50 (+10%)	2.0 (0%)	0.573–0.579
2split*	2.55 (+12%)	1.95 (−3%)	0.896
exp.	2.274	2.009	≈ 0.646

*The 2split C|H walls are free and are held plane, an approximation. The experimental atomisation lies between the two constructions, 2hom the closer (−11%); the C–H bond is right, and the C–C bond is 10% long, as in C_2 (+9%). The carbon model transfers with a consistent C–C bias, plausibly the single σ pair standing in for multiple bonding. Carrying the carbon valence instead as one 2hom pair per atom ($r_c = 0.81$) gives C_2 the right length (+2%) but $D_e = 0.359$ (+55%), a binding of two closed pairs across a free wall that charges nothing, and is not used.

Other linear molecules. The same pseudo-atoms, unchanged (C +2, N +2 with one electron behind, O +3 with a 2hom lone pair behind; each bond one 2hom σ pair, a fixed plane at every heavy nucleus), give the following linear molecules (axisymmetric, $h \approx 0.1$; R_e from a parabola through the scan about $R_{\rm exp}$; atomisation at the observed geometry):

molecule	bond	R_e (exp.)	atomisation (exp.)
HCN	C–N	2.25 (2.179), +3%	0.475 (≈ 0.50)
	C–H	1.97 (2.014), −2%	
CO	C–O	2.19 (2.132), +3%	0.344 (0.413)
CO_2	C–O, both	2.15 (2.192), −2%	0.681 (0.620)
CH	C–H	1.93 (2.116), −9%	0.181 (≈ 0.134)
CH_2, linear	C–H, both	1.97 (2.03, bent)	0.355 (≈ 0.30)

Where carbon sits between two partners on the axis the bonds are within 3%, even with the +3 oxygen of CO and CO_2, whose lone pairs face away from each other, unlike those of O_2. Where carbon ends the chain with an electron behind it (CH, linear CH_2) the C–H bond is short: the remaining carbon electrons of the real radical lie off the axis, which a σ-only model cannot represent.

A chain released. The scans above hold every other bond at its observed length. Released as a whole from the observed geometry, cyanoacetylene HC≡C–C≡N relaxes in all four bonds together (coordinate descent at fixed $h = 0.05$, steps of $0.1\,a_0$, a parabola per bond):

bond	observed (r_s)	relaxed	change
C–H	1.997	1.993	
C≡C	2.273	2.560	
C–C	2.604	2.545	
C≡N	2.190	2.162	

The bonds whose type is set by the atoms (C–H, C≡N) stay within 1.3%; the two carbon–carbon bonds come out nearly equal. The chain construction gives every bond one 2hom pair whatever its order, with the rest of carbon's valence frozen into the +2 core, so a triple and a single C–C bond are the same construction and relax to the same length. Bond alternation, the difference between multiple and single bonds along a chain, is not represented. This is the same defect as the 10%-long C≡C of C_2H_2; the bonds of CO_2, all of one kind, are within 2%.

6.4 Forces at the observed geometry

The force on a nucleus is read as the pull of the converged density, $\mathbf{F}_a = Z_a \int \rho\,(\mathbf{x}-\mathbf{R}_a)/|\mathbf{x}-\mathbf{R}_a|^3$ outside the kernel clamp, plus the nuclear repulsion. It equals $-\partial E/\partial \mathbf{R}_a$ only if the energy is stationary in everything that moves with the nucleus. A bare proton inside a 2hom bond pair carries no wall of its own, and there the readout agrees with the slope of $E(R)$ within 0.002–0.011 (LiH, BeH_2, BH; also CH_4, Section 7). With 2split the H domain's wall moves with H and is placed by continuity, not by the energy, and the readout misses that wall term (BH: −0.040 against +0.086). Forces are therefore read with bond pairs and bare protons; the spread between equivalent H atoms, about ±0.005, is the mesh noise at $h = 0.1$.

For each X–H bond the force on H was read at the observed length and at two shorter ones, 0.1 and $0.2\,a_0$ shorter, all with one construction: every X–H bond one 2hom pair over a bare proton, the heavy atom a pseudo-core with its valence in fixed sectors, $\beta = 0$, $h = 0.1$, angles at their observed values. The model's bond length is where the force on H vanishes, from a quadratic through the three points; k is the stretch stiffness there. Forces in E_h/a_0 (positive: the bond would lengthen), lengths in a_0:

molecule	force at observed	model R_e	observed	error	k
BeH_2 (Be +2)	−0.009	2.476	2.520	−1.8%	0.22
C_2H_2 (C–H)	−0.026	1.963	2.004	−2.1%	0.68
BH (B +3)	−0.017	2.271	2.329	−2.5%	0.31
OH (O +3)	−0.028	1.773	1.832	−3.2%	0.52
C_2H_6 (C–H)	−0.041	1.993	2.067	−3.6%	0.64
C_2H_4 (C–H)	−0.031	1.977	2.054	−3.7%	0.61
CH_4 (C +4)	−0.046	1.974	2.054	−3.9%	0.65
LiH (Li +1)	−0.015	2.886	3.015	−4.3%	0.15
HF (F +3)	−0.041	1.654	1.733	−4.6%	0.61
NH_3 (N +5)	−0.071	1.824	1.913	−4.6%	0.88
H_2O (O +6)	−0.068	1.706	1.809	−5.7%	1.00
NH_3 (N +3)	−0.066	1.795	1.913	−6.2%	0.62

Every X–H bond is short, by 1.8–6.2% (mean about 3.8%), and the error grows across the period, smallest for Be and B and largest for N and O, where the cores carry most valence. Standard

methods do better: Hartree–Fock and DFT under the same constraints are within 2% for every molecule, and RealQM is less accurate than either. The two use different objects. Hartree–Fock's orbitals extend over the whole molecule and overlap everywhere; they are not unique, since any mixing of them among themselves leaves the wave function and every prediction unchanged, so no single orbital has a physical meaning of its own. RealQM's electron densities are local and do not overlap, each a definite charge density in its own region of three-dimensional space, proposed as physical; what they do not yet carry is the cost of pressing such regions together that Hartree–Fock obtains from its orthogonality. The C–H trend among the hydrocarbons comes out in the right order, ethane > ethene > ethyne (1.993 > 1.977 > 1.963; observed 2.067 > 2.054 > 2.004), with about half the observed spread. The stiffnesses are those of bonds that bind too strongly: the CH_4 stretch at $k = 0.65$ puts its symmetric frequency about 20% above the observed one. For the linear molecules of Section 6.3, computed in the axisymmetric reduction, $-\mathrm{d}E/\mathrm{d}R$ over ±10% of the bond is HCN +0.042 (C–N) and −0.013 (C–H), CO +0.045, CO_2 −0.022 per bond, C_2H_2 +0.19 (C–C) and CH −0.046; a force must be read against the stiffness, the triple bond of C_2H_2 carrying +0.19 for a 10% error and CO_2 −0.022 for 2%.

7 Scope, and cases that do not yet work

Scope. Inner shells hold as domains with free walls when the wall is placed by the model's own rule, moved until the densities on both sides are equal: the 1D atoms of the Atom Simulator are built that way (carbon 2+2+2 at −37.879 against −37.845), and an all-electron C_2 in the axisymmetric reduction, with a $1s^2$ core domain around each nucleus inside the 1D atom's core wall, keeps its core and has a C–C minimum (Section 7). They collapse when the wall is moved to lower the energy instead: a free wall charges nothing for the valence pressing against it, so the valence follows the nucleus inward (a 1D carbon whose core wall is chosen by minimum energy squeezes its core to $0.26\,a_0$ and falls to $-61.6\,E_h$), and they cannot be represented in 3D at $h = 0.1$, where a $1s$ shell spans under two cells (every all-electron LiH in 3D collapses, Section 6.1). In 3D the inner shells are therefore excluded by a fixed core radius per atom, set by the atom's ionisation energies; this is a frozen-core model, imposed by the mesh rather than by the model. The axisymmetric reduction carries only pairs symmetric about the axis. Contacts between molecules add a new free wall in a region of high density, and dispersion has no counterpart; intermolecular energies are outside the present model.

Cases that do not yet work. The selection above is a first one. The following were tried beyond it; each is stated with the result found:

- *Heavy atoms carrying their whole valence* (CH_4, NH_3, H_2O): bonds short (CH_4 −5% at $h = 0.07$ and moving shorter with the mesh, about −7% extrapolated); binding two to three times too large, by an amount per heavy atom (ethane twice methane). Carbon is analysed below.

- *Shapes and angles:* the CH_4 square plane lies $0.002\,E_h$ *below* the tetrahedron (CCSD(T) puts it $0.225\,E_h$ above); the H_2O angle closes below 95° with the lone pairs frozen into the core and is pushed open with them as domains of their own.

- *O_2 and F_2:* O_2 over-binds about twofold with a bond 12% short; F_2 has no minimum; no core configuration or radius tried reproduces either.

- *Bond order:* each bond is one pair, so the triple and single C–C bonds of HC≡C–C≡N relax to nearly equal lengths (Section 6.3).
- *Dipoles:* right in sign wherever lone pairs are domains of their own, too large by about two in X–H molecules and half the size in formamide; the two constructions of the bonding pair bracket the observed value (below).
- *Large molecules:* within computational reach, not yet validated (Section 3).
- *Reaction energies:* quotable only where every atom keeps its bonding; an isodesmic test among linear molecules alternates by about ± 2.5 kcal/mol between meshes, below what is needed.

Carbon: what the free wall does not charge. Carbon was taken apart against Hartree–Fock (HF) and CCSD(T), computed with PySCF [7] in the cc-pVQZ basis at the same geometry (C–H $2.054\,a_0$ on tetrahedral directions). RealQM is the 2hom construction on pseudo-C +4 ($r_c = 1.0$), every energy recomputed independently from the converged density. Three measurements point the same way.

Each C–H bond is an independent unit. Adding the n-th hydrogen to CH_{n-1} gains $0.352\,E_h$ in RealQM for $n = 1, \ldots, 4$, constant to 2%; the exact gains grow from 0.13 to 0.19. The atomisation of CH_4 is $1.404\,E_h$ against 0.665. Table 2 splits the fourth bond into kinetic and electrostatic parts. The electrostatic gains agree within 0.05. The difference is the kinetic energy: forming the

Table 2: The fourth C–H bond, $CH_3 + H \to CH_4$: contributions to the binding (E_h; positive binds). Electrostatic: $\Delta(V_{eK} + V_{ee} + V_{KK})$.

	ΔT	electrostatic	ΔE
RealQM, $h = 0.10$	+0.018	+0.329	+0.347
RealQM, $h = 0.07$	+0.013	+0.336	+0.349
HF	−0.143	+0.290	+0.147

bond costs HF $0.14\,E_h$ of kinetic energy, as the virial theorem requires near equilibrium, and costs RealQM nothing. The mesh does not change this: from $h = 0.1$ to 0.07 the H atom's kinetic energy moves from 0.453 to 0.476 (exact 0.5) while the bond's stays near zero.

No shape is preferred. The tetrahedron is preferred by $0.248\,E_h$ in HF and 0.225 in CCSD(T), already in HF and so not a correlation effect. In RealQM, where the bonds do not interact, it is not preferred.

The promotion energy has the wrong sign. At the same r_c, pseudo-C with its four electrons in tetrahedral sectors lies $0.223\,E_h$ *below* the ground-state layout of $2s^2 2p^2$, one 2hom pair and two singles (0.222 at $h = 0.07$). In carbon the four-unpaired state ^5S lies $0.152\,E_h$ *above* ^3P (CCSD(T); HF 0.095). Splitting the pair removes its repulsion (−0.14) and lets it closer to the core (−0.16) for a kinetic cost of 0.08. In carbon the same step puts an electron in a $2p$ orbital, whose angular node costs kinetic energy. The radius r_c, fitted with the tetrahedral layout to the ground-state energy, is therefore fitted to the wrong state, and no choice of layout repairs it.

The common cause. In HF, Pauli exclusion acts through the orthogonality of same-spin orbitals: where two would overlap, one must change sign, and the node costs kinetic energy. This is the repulsion between pairs that spreads them to a tetrahedron and resists contraction, and near the core it is the oscillation that pseudopotentials replace by a repulsion [8]. In RealQM the domains do not overlap, so orthogonality holds without a node: each density ends at a free wall, and the

wall is not charged. The measured consequences are a missing kinetic cost of $0.14\,E_h$ per bond, no shape preference, a promotion energy of the wrong sign, and short bonds. The opposite limit, a density forced to zero at the wall, overcharges. Whether a wall cost between the two exists without a fitted parameter is the main open question of the model. Its test is the CH_4 fourth-bond kinetic cost, the tetrahedron and the promotion energy of carbon recovered, F_2 weakly bound near $2.67\,a_0$, the H_2O angle, and LiH, Li_2, C_2, N_2 left where they are.

A constant wall cost. The simplest candidate is the surface energy $\frac{\beta}{2}\int_{\partial\Omega_k} u_k^2$ of Section 5.3 on every wall, with $\beta = 0.2$. The value is borrowed, not calibrated for this solver or for carbon: in the axisymmetric solver $\beta = 0.21$ gives the exact H_2 binding energy at $R = 2.1$, but at that β H_2's own minimum moves from 1.349 (-3.7%) to 1.507 ($+7.6\%$) and still binds 18% too strongly, so in H_2 a constant β does not give bond and energy together. Here $h = 0.1$, every energy is recomputed independently and the surface term added, with each wall's area taken from its local normal (counting cell faces instead measures a wall exactly only when it lies in a grid plane and over-counts it up to $\sqrt{3}$ when it is tilted; the same ethane turned from the grid axis to the cube diagonal then rose by $0.28\,E_h$, with the normal by 0.001). Ethane at $\beta = 0$ has no C–C minimum: the energy falls, kinetic as well as electrostatic, until the core spheres meet near $2.0\,a_0$, whether the C–C pair is two singles or one 2hom pair. At $\beta = 0.2$ it has one, between 2.3 and 2.9: fits through the four points 2.30–3.30 put it at 2.56–2.73 (-6 to -12%; experiment 2.901). In CH_4 the force on the protons vanishes at 2.047 (experiment 2.054; the energy minimum is 2.01, from 1.97 at $\beta = 0$), the C–H forces in ethane at the observed geometry are below 0.011, and BH at its observed length is left with a force of -0.02. The over-binding is reduced and remains: atomisation $1.25\,E_h$ for CH_4 (from 1.40; exact 0.665), 0.176 for BH (0.135). Restricting β to walls between different atoms, so that carbon's own sector walls stay free, is too weak at the same value: ethane then still shortens. The shape is right in sign only. With the sectors imposed, the tetrahedron lies $0.015\,E_h$ below the square plane (at $\beta = 0$ the square is 0.002 lower; CCSD(T) puts it 0.225 higher), about a fifteenth of the preference. With carbon's walls left free at $\beta = 0$, so that the partition and the forces on the protons decide, CH_4 moved along the path from the tetrahedron to the square plane flattens: the force points toward the square at every angle and the square ends $0.052\,E_h$ lower. Four unit densities on a bare +4 kernel do settle near tetrahedral [20]; four pairs bound to protons do not. A constant β thus gives C–H bond lengths with tetrahedral carbon imposed, but not the C–C bond, the size of the shape preference, the fourth-bond kinetic cost or the promotion energy. The same value that leaves C–C 10% short makes H_2 8% long: the cost a wall must carry differs between the wall inside a bonding pair and the walls between pairs, which is the Pauli cost, and no single constant supplies both.

Carbon–carbon: fixed directions collapse, a free cone holds. With carbon's four valence electrons in fixed tetrahedral sectors, the C–C bond has no minimum: the energy of ethane falls until the core spheres meet near $2.0\,a_0$ (experiment 2.901), and so does bare C_2 built from the same carbon, without hydrogen. Neither a frozen core held at the atom's edge density, nor the hydrogens, nor the size of r_c is the cause: in the axisymmetric reduction an all-electron C_2, with a $1s^2$ core domain around each nucleus (inside the 1D atom's own core wall, 0.700) and four valence electrons in one domain, has a minimum at 1.94, and a rigid core gives the same minimum as one solved with the molecule. Splitting that valence into one bond electron in a fixed cone of 60° toward the partner and three behind brings the collapse back: the bond electron has no kernel in its domain and slides to the midplane at no cost. Letting the cone wall be a free boundary under the model's own rule, moved until the densities meet with equal values, removes it: the cone opens

from 60° to about 67° at the minimum, and C_2 binds at $R_e = 2.13$ ($h = 0.1$) and 2.16 ($h = 0.05$), −8% against 2.348, with no wall cost and nothing fitted. The binding remains about four times too strong. Minimising the energy over the cone angle instead would run away to narrow cones. The 3D sector model imposes its walls; whether free walls under continuity give carbon its shape in 3D is the open test.

Dipoles, and the bracket. At the observed geometry ($h = 0.1$, the constructions of the geometry runs): LiH 6.54 D (experiment 5.88); with lone pairs as domains, H_2O 4.09 (O +6; 1.855), NH_3 3.71 (N +5; 1.47), HF 3.65 (1.83), OH 3.63 (1.66), BH 3.50 (1.27), all with the right sign; with the lone pair frozen into the core, NH_3 inverts (0.46 D, H negative) and H_2O (O +3) has half the size (1.01). Formamide, with every heavy atom in three wedges and the amide π lone pair in the N core, has 1.90 D (experiment 3.73) pointing within about 20° of the O–NH_2 axis. All-electron OH and HF in the axisymmetric reduction (free cone, core from the 1D atom) separate the parts: with the bonding pair over the bare proton, the pair and the proton are nearly neutral and the whole dipole is the lone pairs' displacement behind the heavy atom, 3.4–3.6 D for OH and 3.1–3.9 D for HF; with H keeping its own electron (2split, both walls by continuity) the dipole inverts, −0.72 and −0.44 D, the H electron sitting beyond its proton. The two constructions bracket the observed dipole as they bracket the H_2 energy. Letting charge flow across the split wall until the chemical potentials agree does not select a state between them: the energy falls all the way toward $O^- \cdots H^+$. Polarity is therefore predicted in sign; its size depends on the construction.

8 Conclusion

first selection	status
H_2, He	exact value between 2split and 2hom; H_2 bond −2.3%
X–H bonds (LiH, BeH_2, BH, OH, HF, CH_4, $C_2H_{2,4,6}$, NH_3, H_2O)	from the force on H, all short by 1.8–6.2% (mean 3.8%), most for N and O; residual forces inward, 0.009–0.071 E_h/a_0; C–H order ethane > ethene > ethyne right; LiH dipole right sign
X_2	Li_2, Na_2, B_2, C_2, N_2: bond within 10%, binding between the two constructions
linear molecules	HCN, CO, CO_2 within 3%; C_2H_2 C≡C 10% long
C_2, all-electron	with a free (continuity) cone wall: R_e −8%, binding four times too strong; fixed tetrahedral directions collapse C–C
dipoles	right sign with lone pairs as domains; size between the 2hom (about ×2) and 2split (inverted) constructions
beyond the selection	open questions with results so far (Section 7): carbon's shape in 3D (tetrahedron, angles, torsion), heavy atoms with their whole valence, O_2 and F_2, bond order, reaction energies

Solved as it stands, the hard-partition model of RealQM carries every electron pair in one of two parameter-free constructions, 2split and 2hom, the limits of complete and of no avoidance; for H_2 and He the exact energy lies between them. A shared pair (core, bonding pair) is 2hom, and for LiH and BH the 2hom bond is the closer to experiment. Inner shells hold under the model's own continuity rule for walls, but in 3D at the present mesh they are excluded by a fixed core radius; with that, LiH as one 2hom bonding pair has its bond within 5%, its binding within 3% and a dipole of the right sign. Nothing is adjusted to molecular data; each atom contributes one core radius, set from its ionisation energies.

On this first selection of basic examples the model is a geometry model: bond lengths of hydrides, of X_2 from Li_2 to N_2 and of linear molecules built from the same pseudo-atoms are within a few

per cent, with small residual forces at the observed geometry. Binding energies are a diagnostic. Molecules beyond the selection are not yet reproduced or only with low accuracy, and are left as open questions with the results obtained so far; most point to one of them, whether a repulsion between pairs pressed together is needed. Energies above five domains rest on an independent recomputation from the converged density, which is how a fault in the solver's default electron–electron path was found. The present scope, then: bond lengths a few per cent short across hydrides, X_2 and linear molecules; the C–C bond of C_2 held by the model's own free-boundary rule once carbon's directions are not imposed; polarity right in sign, its size bracketed by the two constructions of the bonding pair; shape (angles, torsions) imposed, not predicted; energies a diagnostic. The validation continues: the selection will be extended molecule by molecule, and the cases that do not yet work are the next steps. The purpose throughout is not to compete with standard quantum chemistry in accuracy but to establish what a computable three-dimensional continuum model of the atomic world achieves as it stands, and where and why it falls short. The failures reported here are as informative as the successes: each exposed a specific feature of the model, and together they point to one missing ingredient, the cost of pressing separate electron groups together, and to the continuity rule for walls as the part that already works. In detail, each with the step it suggests:

- *C–C collapse* (ethane, bare C_2, and C_2 with a frozen core): imposed tetrahedral directions leave the bond electron in a domain with no kernel and nothing behind it. Solved for C_2: with the cone wall free under the continuity rule the bond holds (-8%). Next: the same rule for the sector walls in three dimensions.

- *Shells collapsing when a wall is chosen by energy* (1D carbon): energy minimisation over a free wall lets the valence follow the nucleus inward. Solved: the continuity rule, which the Atom Simulator already uses.

- *No shape* (CH_4 flattens, propane folds, no twist barrier): the walls between a heavy atom's own electron groups cost nothing. Next: continuity for those walls in three dimensions, as above; if that is not enough, a cost charged when separate groups press together, growing with the density at the wall rather than constant.

- *X–H bonds short by* 2–6% where Hartree–Fock is within 2%: the same missing cost, measured, largest for N and O. The step is the same as for shape.

- *The brackets* (H_2 energy between 2hom and 2split; OH and HF dipoles too large with 2hom and inverted with 2split): what lies between is the correlation of the two electrons of a pair. No rule within the model selects it: an overlap penalty and a charge balance across the wall were tried and both run to an end. Open; until a rule is found the bracket is the stated uncertainty.

Criteria of the model's own. Excused from competing in accuracy, a model must state what it can fail. RealQM sets itself these, scored on the present results:

criterion	present status
bond lengths from the forces, nothing fitted, within a few per cent	met: X–H 2–6% short, C_2 −8%
shell structure of atoms from the continuity rule	met: $Z \leq 18$ within about 2%, carbon 0.09%
sign of polarity	met where lone pairs are domains of their own
molecular shape (tetrahedron, angles, torsion) from the model	not met: imposed
energies within a stated bracket	met as a bracket, not closed
computable at the scale of proteins	estimated, not demonstrated

Understanding. Standard quantum mechanics offers understanding through models built on it (orbitals, hybridisation, the Pauli principle, frontier orbitals) whose objects are constructions and not unique. RealQM offers understanding of another kind: a picture that is the model itself, electrons as charge densities in ordinary space, a bond as charge between nuclei, exclusion as non-overlap, shells from the continuity rule, and one that shows where such a picture succeeds and where it fails. A picture is understanding only as far as it is right: where RealQM fails, at shape and at the energies, its account is incomplete, and the failures say where.

Moderate accuracy at low cost. H_2 shows the profile that may carry over to large molecules: the geometry within a few per cent, the energy bracketed by the two constructions of the pair. In accuracy RealQM does not compete with the computable methods of standard quantum chemistry; where it could compete is cost. With each domain kept on a local box and the potentials solved by multigrid, one iteration touches every grid cell once, so the cost is set by the number of grid points alone, (volume)$/h^3$, linear in the size of the molecule and independent of the number of electrons beyond the volume they fill. On one GPU at about 10^{10} cell updates per second, this estimates a peptide of about 100 atoms in seconds, 10^3 atoms in about a minute, and a protein of 10^4 atoms (10^8–10^9 cells at $h = 0.2$–0.1) in minutes to an hour. In practice DFT is a family of approximate functionals, calibrated or validated against classes of data and applied to similar systems; how well it works is known where it has been checked, not in general. Kohn–Sham DFT grows as N^3 and is routine to a few thousand atoms (a thousand atoms takes hours on hundreds of cores, or about an hour on a few GPUs); linear-scaling DFT reaches 10^4–10^5 atoms in insulators at about 10^3–10^4 core-hours for 10^4 atoms. Allowing for the throughput of a GPU against CPU cores, the estimate leaves RealQM one to two orders of magnitude cheaper at equal size, from its short stencil on a uniform grid with no basis and no matrix algebra. These are estimates for a solver not yet written in that form. The present `molecule.js` holds one amplitude field with each cell labelled by its domain, but its validated electron–electron path solves a Poisson problem over the whole grid for each electron and keeps each result, so that part grows about as (domains)×(grid); a local path, with one total potential and each domain's self-potential on a small box around it, is in development. With explicit imaginary-time steps a small molecule takes minutes. The aim they describe is moderate accuracy at low cost: geometries within a few per cent for systems where DFT becomes expensive.

References

[1] C. Johnson. The RealQM programme is set out in three submitted articles: *Quantum Mechanics as Multiphase 3D Continuum Mechanics* (RealQM), submitted to the *Annales de la*

Fondation Louis de Broglie; *A 3D Multiphase Continuum Computational Model for Atoms*, submitted to the *International Journal of Quantum Chemistry*; and *RealNucleus: Nuclear Binding as Dual Coulomb Confinement, without a Strong or Weak Force*, submitted to *Physics Essays*.

[2] C. Johnson, *RealQM for Atoms and the Periodic Table*. The shell configurations used here, including oxygen as 2+3+3, are obtained there from atomic calculations independent of any molecular input.

[3] K. P. Huber and G. Herzberg, *Molecular Spectra and Molecular Structure IV: Constants of Diatomic Molecules*, Van Nostrand Reinhold (1979).

[4] NIST Computational Chemistry Comparison and Benchmark Database, NIST Standard Reference Database 101.

[5] Public code, the molecule simulator, and an interactive gallery: `https://claes542.github.io/RealMolecule/gallery.html` (source repository: `https://github.com/Claes542/RealMolecule`).

[6] D. Vasiliu, K. A. Peterson and D. A. Dixon, Benchmark-quality atomization energies for BeH and BeH_2, *J. Chem. Theory Comput.* **13**, 649 (2017).

[7] Q. Sun *et al.*, Recent developments in the PySCF program package, *J. Chem. Phys.* **153**, 024109 (2020).

[8] J. C. Phillips and L. Kleinman, New method for calculating wave functions in crystals and molecules, *Phys. Rev.* **116**, 287 (1959).

[9] J. C. Slater, The virial and molecular structure, *J. Chem. Phys.* **1**, 687 (1933).

[10] R. P. Feynman, Forces in molecules, *Phys. Rev.* **56**, 340 (1939).

[11] H. Hellmann, Zur Rolle der kinetischen Elektronenenergie für die zwischenatomaren Kräfte, *Z. Phys.* **85**, 180 (1933).

[12] K. Ruedenberg, The physical nature of the chemical bond, *Rev. Mod. Phys.* **34**, 326 (1962).

[13] G. B. Bacskay and S. Nordholm, Covalent bonding: the fundamental role of the kinetic energy, *J. Phys. Chem. A* **117**, 7946 (2013).

[14] W. Kołos and L. Wolniewicz, Potential-energy curves for the $X\,^1\Sigma_g^+$, $b\,^3\Sigma_u^+$, and $C\,^1\Pi_u$ states of the hydrogen molecule, J. Chem. Phys. **43**, 2429 (1965).

[15] W. Kołos and C. C. J. Roothaan, Accurate electronic wave functions for the H_2 molecule, *Rev. Mod. Phys.* **32**, 219 (1960).

[16] W. Kołos and L. Wolniewicz, Improved theoretical ground-state energy of the hydrogen molecule, *J. Chem. Phys.* **49**, 404 (1968).

[17] G. Herzberg, The dissociation energy of the hydrogen molecule, *J. Mol. Spectrosc.* **33**, 147 (1970).

[18] C. A. Coulson and I. Fischer, Notes on the molecular orbital treatment of the hydrogen molecule, *Phil. Mag.* **40**, 386 (1949).

[19] M. Conti, S. Terracini and G. Verzini, Asymptotic estimates for the spatial segregation of competitive systems, *Adv. Math.* **195**, 524 (2005).

[20] C. Johnson, *The Tetrahedral Stereocentre from Packing: Chirality's Geometric Origin without Hybridisation*, preprint (2026).