

Upscaling of RealQM towards Hilbert's Sixth Problem

Claes Johnson¹

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

October 16, 2026

Abstract

Hilbert's sixth problem asks, in the part Hilbert attributed to Boltzmann, for a rigorous passage from the atomistic description of matter to the equations of continuous media. The classical route — Newtonian particles, then Boltzmann's kinetic equation, then Euler and Navier–Stokes — was carried through for hard spheres by Deng, Hani and Ma in 2025. That achievement comes with a structural restriction that is not a defect of the proof but of the route: the Boltzmann–Grad scaling sends the volume fraction to zero, so what is reached is a *dilute gas*. Liquids, solids and every dense phase lie outside it by construction. This programme proposes the complementary route. RealQM has no particles to coarse-grain: it is already a continuum theory, N non-overlapping charge densities on ordinary three-dimensional space, with the interfaces free. RealQM's continuum lives on the atomic scale and continuum mechanics on the laboratory scale, so the two descriptions must still be connected — but for the *equilibrium* coefficients no averaging over microscales is needed at all. An elastic modulus is a second derivative of the ground-state energy with respect to a macroscopic deformation, the deformation can be imposed exactly on a periodic cell as a boundary condition, and the energy is then one number computed exactly. No field is averaged; a scalar is differentiated. The whole of linear elasticity, the equation of state, the force constants and hence the phonons are reachable this way, and sodium's equation of state has already been obtained so — r_s and a bulk modulus from Coulomb geometry, with no exchange and no band-filling term. We set out what is already exact, what this direct route reaches, and what it does not. Exact and by the chain rule alone: RealQM's localisation energy *is* the von Weizsäcker functional, its variation *is* Bohm's quantum potential, and the theory as used is precisely the zero-velocity equilibrium of a Madelung pair — continuity plus a momentum balance whose stress is that potential. That balance has now been verified numerically at *two* domains, where no Schrödinger equivalent stands behind it, and the measurement locates the accuracy-limiting feature of the construction at the free interface — fourth-order convergence in the bulk against first order there, which is the surface on which dissipation is being sought. To be supplied: the dynamics of that pair per domain, with the interfaces moving as free surfaces. And the direct route stops at a sharp boundary, which is the useful thing it tells us: it reaches every coefficient that is a derivative of an energy, and no coefficient that is not. Genuinely open, and now the *sole* open problem: **dissipation**. Madelung–Euler is non-dissipative, Navier–Stokes is not, and RealQM has no statistical ensemble from which an H -theorem could extract irreversibility. Where viscosity comes from, if it comes at all, is the question this project exists to answer. We state it as a question rather than announce an answer.

1 What Hilbert asked, and what has been delivered

The sixth of Hilbert’s 1900 problems is headed *Mathematical Treatment of the Axioms of Physics*. Its second paragraph is the part at issue here: Hilbert points to Boltzmann’s work on the limiting processes “which lead from the atomistic view to the laws of motion of continua” and asks for them to be made rigorous. The programme it names is therefore specific — not an axiomatisation of physics in general, but a derivation of the equations of continuum mechanics from a description of matter in terms of its constituents.

The classical route runs in two stages. From Newtonian dynamics of many particles to Boltzmann’s kinetic equation, in the Boltzmann–Grad scaling; then from Boltzmann to the fluid equations, in a hydrodynamic scaling. The first stage is where the difficulty has always been concentrated, because it must produce irreversible macroscopic behaviour from reversible microscopic laws and must do so for times long enough to matter. Deng, Hani and Ma completed it for hard spheres on the two- and three-dimensional torus, obtaining the compressible Euler and incompressible Navier–Stokes–Fourier equations with quantitative error bounds [1]. It is reasonable to call this a solution of the problem for gas dynamics.

The restriction, which is the opening for this programme. The Boltzmann–Grad limit takes N particles of diameter ε with $N\varepsilon^2$ fixed, so the volume fraction $N\varepsilon^3 \rightarrow 0$. The matter described is therefore arbitrarily dilute, and the critical literature has made the point sharply: the result characterises a dilute gas and cannot embody the properties of a dense fluid [2]. This is not a gap in the argument to be repaired. It is what the route is: kinetic theory is a theory of free flight interrupted by rare binary collisions, and a liquid is not that.

So the honest statement of where the problem stands is narrower than the headline. Continuum mechanics has been derived from atomism *for gases*. For the condensed phases — the ones chemistry, materials science and biology are made of — it has not, and the classical route does not reach them.

2 The inversion

RealQM [4, 5] describes an N -electron system by N non-negative densities ψ_i , each carrying unit charge on its own domain Ω_i , the domains non-overlapping and their interfaces free to move to where the energy is least. The energy is Coulomb’s:

$$E[\{\psi_i\}, \{\Omega_i\}] = \sum_i \int_{\Omega_i} \left(\frac{1}{2} |\nabla \psi_i|^2 - \sum_a \frac{Z_a}{|x - x_a|} \psi_i^2 \right) dx + \frac{1}{2} \sum_{i \neq j} \iint \frac{\psi_i^2(x) \psi_j^2(y)}{|x - y|} dx dy, \quad (1)$$

subject to $\int_{\Omega_i} \psi_i^2 = 1$.

The observation that organises this programme is that (1) is not an atomistic description at all. There are no particles in it. There are continuous fields on ordinary three-dimensional space, and the only discreteness is in the *number of domains* — an integer labelling of regions, not a collection of point masses with positions and momenta.

The consequence for Hilbert’s problem is structural, but it must be stated carefully, because the tempting version of it is wrong. RealQM is a continuum theory *at the atomic scale*: its features vary over lengths of order a_0 . Macroscopic continuum mechanics describes variation over laboratory lengths L . So an averaging step is still required — the two descriptions are not

the same description — and it would be an overclaim to say that no passage between scales is needed.

What changes is *which* passage. The classical route must build a continuum out of particles, and every difficulty it faces — the scaling limit, the propagation of chaos, the recovery of irreversibility — belongs to that construction. The route here passes from a continuum on scale a_0 to a continuum on scale L , with small parameter

$$\varepsilon = a_0/L, \quad (2)$$

which is *not* a dilution parameter and carries no requirement that matter be rare.

The formal apparatus for such a passage is homogenisation — two-scale convergence, G - and H -convergence, periodic unfolding — and it is available here in a way it is not available to the standard formalism (Section 9). But for the coefficients of equilibrium continuum mechanics **it is not needed**, and Section 3 is the argument that it should not be used. They can be computed directly, without averaging anything over the microscale, and one of them already has been.

And because no dilution is involved, the question is being asked exactly where the classical route cannot go. RealQM’s established regime is atoms, molecules, condensed phases and proteins [5]. If continuum mechanics follows from (1), it follows for dense matter.

3 Direct upscaling: no averaging required

Consider what an elastic modulus actually is. It is a second derivative of the ground-state energy with respect to a macroscopic deformation,

$$C_{ijkl} = \frac{1}{V} \frac{\partial^2 E}{\partial \epsilon_{ij} \partial \epsilon_{kl}}, \quad B = V \frac{\partial^2 E}{\partial V^2}, \quad (3)$$

and the deformation ϵ is a *macroscopic* quantity that can be imposed on a periodic cell exactly, as a boundary condition: one changes the shape or the volume of the cell. RealQM then computes the ground state in that cell and returns a single number, its energy. The modulus is the curvature of that number as a function of the imposed deformation.

Nothing has been averaged over the microscale. No stress field is integrated, no two-scale expansion is posed, no corrector problem is solved. A scalar is computed exactly and differentiated with respect to a macroscopic parameter. The only use made of the cell is that the macroscopic state is uniform across it — which is the content of $\varepsilon \ll 1$ in (2), carried as an assumption rather than as an asymptotic analysis.

What this reaches. Every coefficient of continuum mechanics that is a derivative of the ground-state energy with respect to a macroscopic control parameter:

- the full elastic tensor and the bulk, shear and Young moduli;
- the equation of state $P(V) = -\partial E/\partial V$, and with it the speed of sound;
- the force constants $\partial^2 E/\partial u_a \partial u_b$ and hence the phonon spectrum, and through it the thermal expansion in the quasi-harmonic treatment;
- the dielectric and piezoelectric response, as derivatives with respect to an applied field;
- surface and interface energies, as differences of computed energies.

That is the whole of equilibrium continuum mechanics. It requires no new machinery — only the ground-state solver the programme already has.

A worked example that already exists. Sodium’s equation of state has been computed this way on a Wigner–Seitz cell, one domain per ion filling the cell with the normal derivative vanishing at the face by symmetry — which is RealQM’s own prescription and coincides with a construction in use since 1933. The energy curve $E(r_s)$ is minimised for the equilibrium density and differentiated twice for the modulus:

	RealQM	standard	measured
equilibrium r_s	$3.77 a_0$	$3.87 a_0$	$3.93 a_0$
bulk modulus	9.1 GPa	8.0 GPa	6.3 GPa

The RealQM column keeps the same electrostatics as the standard one and differs only in the kinetic energy — the cell state’s own $\langle T \rangle$ in place of a band-filling term, and no exchange. A macroscopic material constant, from Coulomb geometry, with nothing fitted.

The accuracy should be read with the same care as everywhere else in this programme: 9.1 against 6.3 GPa is 44% high, and worse than the standard column’s 27%. A bulk modulus is a second derivative, so it amplifies the model’s error in the way H_2 ’s binding energy amplifies the error on its total. The point of the example is not that the number is good. It is that the number exists, was obtained without averaging, and required no ingredient the theory did not already have.

And where it stops, which is the informative part. The direct route reaches a coefficient if and only if that coefficient is a derivative of an energy. *Viscosity is not.* A viscous stress is dissipative: the work done against it is not stored, so there is no $E(\text{strain rate})$ to differentiate, and no variational principle whose minimiser contains η . This is exactly why the standard treatment reaches for Green–Kubo, which extracts a transport coefficient from equilibrium fluctuations and therefore needs an ensemble.

So the division of the problem is sharp, and it is not an artefact of the method chosen:

	conservative coefficients	dissipative coefficients
examples	elastic tensor, $P(V)$, phonons	viscosity, thermal conductivity
nature	derivative of ground-state energy	transport, no energy functional
this route	direct, no averaging	not reached

Equilibrium continuum mechanics follows from RealQM by differentiating computed energies. Dissipation does not follow at all, and is the whole of what remains open (Section 8).

What is given up, stated plainly. The direct route delivers the *coefficients* of the macroscopic laws, not their *form*. Writing (3) presupposes that E is quadratic in the strain near equilibrium — that is, it assumes Hooke’s law rather than deriving it. Homogenisation would instead produce linear elasticity as the macroscopic limit, form and coefficient together, and that is a genuine advantage of the formal route which should not be waved away.

Two things reduce the force of the objection. The form can be established computationally: compute $E(\epsilon)$ across a range and verify that it is quadratic, and that the departure from quadratic sets in where nonlinear elasticity is observed to set in. That is a demonstration and not a theorem. And a demonstration is the kind of evidence this programme is built to give — its standing claim is that its assertions are testable *because* computable, and reaching instead for

an asymptotic theorem would borrow precisely the sort of authority it declines to grant to the standard formalism elsewhere. The formal route remains worth having; it is not what the next computation should wait for.

4 What is already exact

The first part of the bridge is built and requires no approximation and no limit; it is reported in full in [3] and summarised here because the programme starts from it.

The localisation energy is the von Weizsäcker functional. Writing $\rho_i = \psi_i^2$ with $\psi_i \geq 0$, the chain rule gives $\nabla \rho_i = 2\psi_i \nabla \psi_i$, hence

$$\frac{1}{2} \int |\nabla \psi_i|^2 dx = \frac{1}{8} \int \frac{|\nabla \rho_i|^2}{\rho_i} dx = T_{\text{vW}}[\rho_i], \quad (4)$$

in units $\hbar = m = 1$. An identity, not a model.

Its variation is the quantum potential. With $u = \sqrt{\rho}$,

$$\frac{\delta T_{\text{vW}}}{\delta \rho} = -\frac{1}{2} \frac{\nabla^2 \sqrt{\rho}}{\sqrt{\rho}} = Q, \quad (5)$$

Bohm's quantum potential — obtained here not by analogy but as the variation of a term (1) writes down for an unrelated reason, namely the cost of confining a charge to a region.

The theory as used is a hydrostatics. Stationarity of (1) is the statement that $V + Q$ is constant on each domain, which is the $v = 0$ case of the Madelung pair

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho v) = 0, \quad (6)$$

$$\frac{\partial v}{\partial t} + (v \cdot \nabla)v = -\nabla(V + Q). \quad (7)$$

Equation (6) is conservation of charge and (7) is a momentum balance whose stress is supplied by (5). *RealQM is therefore already the equilibrium of a continuum mechanics, and the constitutive law is not postulated — it is the variation of the energy the theory was written with.* This is the sense in which the derivation asked for by Hilbert is, on this route, partly a matter of recognition rather than of proof.

And the interfaces carry a surface tension. The free boundary between two domains contributes an energy per unit area, so the multiphase structure is a continuum mechanics *with surfaces*, which is what dense matter requires and what a kinetic theory has no room for [3].

5 Mass distribution and forces

A theory of continuum mechanics needs a mass density and a stress. RealQM as written supplies neither directly — it is a theory of *charge* — so the two must be identified before any averaging can be done. Both identifications already exist in this programme, which is worth saying because it means the present project inherits them rather than inventing them.

Mass is not a separate postulate. In the account given in [7], mass is not primitive: it is energy showing up as inertia, through the measured law $\mathbf{p} = (E/c^2)\mathbf{v}$. The mass density is therefore

$$\rho_m = \frac{\text{energy density}}{c^2}, \quad (8)$$

and the energy density is explicit in (1). Nothing is added; the mass field is read off the energy field.

The stress is the localisation term. The same work makes the stress tensor explicit, identifying the localisation energy — the term that became T_{vW} in Section 4 — as a cohesive stress, and the stationary Schrödinger equation as the statement that this stress is in local balance [7]. So the Cauchy stress of the macroscopic theory is to be obtained by averaging an object the microscopic theory already possesses, not by fitting a constitutive law.

But the mass is not where the charge is, and this splits the problem in two. The electrons carry the charge and essentially none of the mass; the kernels carry essentially all of it, some 99.95%, and in the Born–Oppenheimer treatment used throughout they are point charges at fixed sites. So:

- the charge and the stress are continuous on scale a_0 , and their passage to the macroscopic scale is **continuum to continuum** — homogenisation, as in Section 2;
- the mass sits at **discrete** sites, and its passage is a lattice average, the classical atomistic-to-continuum problem for a crystal — well understood, and not the hard part, because the kernels are slow and nearly periodic.

The macroscopic momentum balance to be derived is then

$$\rho_m \left(\frac{\partial v}{\partial t} + (v \cdot \nabla)v \right) = \nabla \cdot \sigma, \quad (9)$$

with ρ_m from the lattice average of (8) and σ from the homogenised localisation stress.

And this is where dissipation plausibly lives. The split just made is also the physical mechanism the programme was missing. The electron densities are fast and relax essentially instantaneously; the kernels are slow and massive. Work done on the free boundaries while the kernels move is not returned to the kernels, and that transfer — electron–phonon coupling, in the usual language — is a dissipation mechanism internal to the model rather than an extra postulate. It subsumes the “interface motion” candidate of Section 8 and gives it a carrier.

6 Two objections that narrow the claim

Two criticisms of the programme as stated above are correct and change what may be claimed. They are recorded here rather than left for a reader to supply.

A charge continuity equation is not a mechanical momentum balance. The Madelung pair (6)–(7) governs the flow of *charge*: ρ is a charge density and ρv a charge current. The momentum of continuum mechanics is $\rho_m v$ with ρ_m from (8), and these are different objects — the first is conserved by (6), the second is what (9) requires. Writing (7) does not by itself establish that mechanical momentum is conserved, and three things must be shown before it does:

- that the inter-domain Coulomb forces obey Newton’s third law domain by domain, so that the total of $\rho_m v$ over a closed system is constant;
- that the *electromagnetic field momentum* is accounted for. RealQM’s $i \neq j$ coupling in (1) is instantaneous, so it carries no retardation and no field momentum, and a momentum balance that ignores a momentum reservoir is incomplete. The programme’s own treatment of inertia is explicitly electromagnetic [7], which is where this has to be settled;
- that the identification $\rho_m = E/c^2$ survives averaging, i.e. that the homogenised mass density is the average of the microscopic one and not merely proportional to it.

Until these are in place, Section 4 should be read as establishing a *charge*-hydrodynamics, with the mechanical reading a further step.

Deriving continuum mechanics from microscopic quantum mechanics is not new.

This is the more important correction, because the framing in Section 1 invites the wrong contrast. Standard quantum mechanics already supports microscopic routes to every macroscopic quantity at issue: elastic constants from density-functional calculations, phonons from lattice dynamics, viscosity and thermal conductivity from Green–Kubo relations and from Boltzmann transport. These are used daily and they work for *dense* matter — the restriction that confines Deng, Hani and Ma to a dilute gas is a restriction on their *rigour*, not on practice. So the claim “continuum mechanics has not been derived for condensed phases” is false as a statement about physics, and true only as a statement about theorems.

What remains distinctive is therefore narrower, and should be stated as follows rather than as a claim of priority:

- the starting point is *already* a three-dimensional continuum, so the constitutive law is obtained by homogenising a partial differential equation rather than by statistical mechanics or linear response — a different kind of derivation, with a different kind of error estimate available;
- no exchange–correlation functional and no $3N$ -dimensional wavefunction enter, so there is no fitted ingredient between the microscopic law and the macroscopic coefficient;
- the route is in principle open to *proof*, since homogenisation theorems exist for problems of this type, whereas a density-functional computation of an elastic constant is a recipe and not a theorem.

The third is a hope and not a result, and the first two are claims about the *character* of the derivation rather than about its reach. Nothing here says the macroscopic answers will be better than those standard methods already give — Section 10.

7 What must be supplied

The dynamics. Equations (6)–(7) must be posed per domain, with the interfaces evolving as free surfaces: the kinematic condition that a boundary moves with the normal velocity of the fluid, and continuity of the chemical potential across it. None of this is new physics; it is the completion of a pair whose static half is already in use.

Two qualifications belong with it, and they are stated here rather than discovered later. Madelung hydrodynamics is equivalent to the Schrödinger equation only for irrotational v with quantised circulation, $\oint v \cdot dl = 2\pi n$; both are extra conditions rather than consequences, and they are where the quantum character re-enters a formulation that otherwise reads as fluid mechanics. And the free-surface problem for two Madelung fluids in contact is a genuine free-boundary problem whose well-posedness is not established.

8 What is genuinely open: dissipation

Here is the hard part, and the reason this is a programme and not a result.

Equations (6)–(7) are an *Euler* system. They have no viscosity, no heat conduction and no entropy production. Navier–Stokes has all three. On the classical route the irreversibility is manufactured in the first stage: molecular chaos breaks the time symmetry of the particle dynamics, Boltzmann’s H -theorem supplies an entropy, and the transport coefficients come out of the collision operator as integrals over a scattering cross-section. Every one of those ingredients is statistical, and RealQM has none of them. There is no ensemble, no distribution function, no collision, and no averaging — which is precisely why the first stage was free.

The bill arrives here. **The step from Euler to Navier–Stokes, trivial on the classical route and the whole content of the transport coefficients, has no mechanism in this framework as it stands.** Three candidate sources are worth examining, and distinguishing them is the first real task of the project:

1. **Interface motion as the dissipative degree of freedom.** A system of N electrons has $O(N)$ free boundaries. Their motion is not a small correction to the bulk; in this framework conduction *is* boundary motion [6]. Coarse-graining over many interfaces could produce an effective stress with a dissipative part, in the way that grain-boundary motion produces macroscopic plasticity. Prediction to test: the resulting viscosity should scale with interface area per unit volume, hence with $N^{1/3}$ per domain rather than with a mean free path.
2. **Exchange of territory as an irreversible process.** The partition $\{\Omega_i\}$ is itself a macroscopic variable. Rearrangements of which electron owns which region are discrete events with no reverse-path symmetry once the geometry has relaxed, which is a candidate for entropy production that is geometric rather than statistical.
3. **Finite precision.** The programme’s own thermodynamics locates irreversibility in finite resolution rather than in statistics [8], and the Clay work treats well-posedness as a property of computed output [9]. If that is the right account, viscosity in this framework is a residual of unresolved scales and the Euler system is the only exact one — a conclusion that would be consistent, but which must be argued, not assumed.

Candidate 3 would make the Euler–Navier–Stokes gap a feature; candidates 1 and 2 would close it. They make different predictions and are therefore separable, which is the only reason the question is tractable.

And the target is sharper than “find a dissipation mechanism”. Any of the three can plausibly transfer energy into motions the macroscopic description does not resolve. That is not enough. Newtonian viscosity is a very particular constitutive form: the dissipative stress is *local* in space, and *linear in the strain rate*. So the derivation owes two things beyond the existence of a mechanism — why the resulting stress depends only on the local velocity gradient rather than on the history or the surroundings, and why it is proportional to that gradient rather than to some other function of it. A mechanism that produced a nonlocal or nonlinear stress would not give Navier–Stokes; it would give something else, which might be interesting but would not answer Hilbert. This is the form in which the question should be posed, and the scaling test in candidate 1 is only the first of the two parts.

9 Then what is the difference from standard quantum mechanics?

Section 6 concedes that elasticity, phonons, viscosity and thermal conductivity are all already obtained from microscopic quantum mechanics for dense matter. The obvious question follows: if the macroscopic answers exist, what does this route add? Four differences survive, none of them a difference in accuracy.

1. In standard quantum mechanics, the object homogenised is never the object claimed to be exact. The many-electron Schrödinger equation cannot be homogenised at all: Ψ lives on $\mathbb{R}^{3N}$, and configuration space contains no atomic unit cell over which to average. Every standard route therefore replaces it first — by Kohn–Sham, a three-dimensional equation carrying a fitted exchange–correlation functional; or by a parameterised potential-energy surface; or by correlation functions supplied to a Green–Kubo formula. The chain reads

postulate $\longrightarrow$ uncontrolled approximation $\longrightarrow$ 3D surrogate $\longrightarrow$ macroscopic law,

and no error estimate crosses the first arrow, because the exact object was never evaluated. The route proposed here has

postulate $\longrightarrow$ homogenisation $\longrightarrow$ macroscopic law,

with a single approximation step carrying a controlled small parameter $\varepsilon = a_0/L$. This is the structural difference, and it is the only one of the four that is a difference in kind.

2. The cell problem is posable. Homogenisation requires the microscopic solution on one unit cell under periodic conditions. In RealQM that is an N -electron three-dimensional solve at cost linear in N . Posed honestly in standard quantum mechanics it is a $3N$ -dimensional solve on the cell, which is out of reach beyond a few electrons — which is precisely *why* the standard routes use surrogates. The surrogates are not a convenience; they are forced.

3. Dissipation would be a different kind of explanation, not a different number. Green–Kubo obtains viscosity from an ensemble in a thermal state through the fluctuation–dissipation theorem: the derivation is irreducibly statistical. RealQM has no ensemble, so if a viscous stress emerges here it must be geometric — free boundaries moving against slow kernels, as in Section 5. The coefficient might agree; the account of why there is a coefficient would not.

4. No exchange–correlation functional. A density-functional elastic constant depends on which of several hundred fitted functionals is chosen, so it is not parameter-free. One obtained here would be, subject to the model’s own error.

What is not a difference. Accuracy is not: a density-functional calculation will give a better elastic constant than this will, and Section 10 says so. Deriving continuum behaviour from microscopic physics is not: that is daily practice. Phonons are not: they need the same Born–Oppenheimer machinery here as there, with no advantage.

The discriminating test, and it is not an elastic constant. If this route reproduces the same macroscopic equations with poorer coefficients, the differences above are epistemic, and a reader is entitled to ask whether they matter. There is one place where the difference is not epistemic at all. **The statistical route is not merely difficult where local thermal equilibrium fails — it is invalid.** Green–Kubo presupposes an ensemble in local equilibrium, and that presupposition fails for strongly confined flow, for very high strain rates, for glasses, and for non-equilibrium steady states. A continuum theory obtained by homogenising a deterministic field theory carries no such requirement and still applies.

So the test to aim at is not a coefficient that density-functional theory computes better. It is either the *functional form* of the viscosity — the interface-area scaling of Section 8 against a mean-free-path dependence, which differ — or a regime in which no local equilibrium exists, where the standard derivation has nothing to say and this one would still deliver an equation. That is a claim about scope rather than precision, which is the form in which this programme’s claims have been most defensible.

10 A note on what this cannot claim

Since Section 6 concedes that the macroscopic quantities are already obtainable by established methods, the programme’s value cannot rest on reaching them. It rests on the character of the passage: no fitted functional, no high-dimensional wavefunction, a constitutive law from a cell problem, and — if it can be had — a theorem. Against that, the model’s own accuracy is a few per cent, so no macroscopic coefficient computed this way should be expected to beat a density-functional value. If that is the test, this programme loses it. The test it can win is whether a single parameter-free three-dimensional continuum theory yields the laws of continuous media without a statistical limit.

11 Phase 1, done: the identity verified, and a warning about the route

The first item of the programme has been carried out, and it returned one result and one caution.

The caution, which redirected the work. Integrating the Madelung pair (6)–(7) directly is a harder numerical problem than it looks, and for a structural reason: Q contains $\nabla^2 \sqrt{\rho}$, so ∇Q in the momentum equation carries a *third* derivative of the density. The pair is therefore dispersive, in the manner of the Korteweg–de Vries equation, and explicit time-stepping is restricted to $\Delta t \lesssim C\Delta x^3$ rather than the $\Delta t \lesssim C\Delta x$ one would budget for. This is a fact about the formulation and not about any particular code. It is also the reason the “missing half” has not simply been implemented already.

The result. For a *single* domain the Madelung pair is equivalent to the Schrödinger equation, so a bespoke Madelung integrator would buy no new physics. The informative test is therefore whether the hydrodynamic quantities read off an accurately computed Schrödinger solution satisfy (6)–(7). Taking a coherent state of a harmonic trap, propagated spectrally, and computing every time derivative analytically from the Schrödinger right-hand side so that no time-differencing error enters, the two residuals

$$R_1 = \rho_t + (\rho v)_x, \quad R_2 = v_t + vv_x + (V + Q)_x$$

vanish under refinement at fourth order — the order of the difference operators used to evaluate them. Relative to the magnitude of the terms they cancel between, R_1 reaches 1.2×10^{-8} and R_2 reaches 1.1×10^{-7} ; the floor on R_2 is round-off in $\nabla^2 \sqrt{\rho} / \sqrt{\rho}$ where ρ is small, and it moves with the density threshold exactly as that explanation requires, so it is a property of the diagnostic rather than of the identity.

R_1 is probability conservation and was never in doubt. R_2 is the claim that matters: that the quantum potential *is* the stress of a continuum mechanics. In one dimension, for one domain, it holds to seven figures. *This verifies a claim the programme had already asserted; it does not extend it.* The new content begins at two domains, where there is no Schrödinger equivalent to check against.

12 Phase 2, done: the momentum balance is an identity at two domains

Programme item 2 below has been carried out, and it passes. It is the first test with two domains and a free interface between them, and — unlike Phase 1 — it has no Schrödinger equivalent standing behind it.

The stress, written down. With $u = \sqrt{\rho}$ and $Q = -\frac{1}{2}\nabla^2 u/u$, the chain rule alone gives

$$\rho \nabla Q = \nabla \cdot T, \quad T_{ij} = -\frac{1}{2} [u \partial_i \partial_j u - \partial_i u \partial_j u] = -\frac{1}{4} \partial_i \partial_j \rho + \frac{\partial_i \rho \partial_j \rho}{4\rho}, \quad (10)$$

with no δ_{ij} term; in one dimension this is $\sigma = \frac{1}{4}[(\rho')^2/\rho - \rho'']$. The static Euler–Lagrange equation of the model is $Q + V = \lambda$ with $V = -K + P_{\text{other}}$, so the $v = 0$ momentum balance is the statement

$$\nabla \cdot T = -\rho \nabla V, \quad (11)$$

and integrating its z -component over a half space $\{z > z_{\text{cut}}\}$ turns it into a statement about forces: the flux of quantum stress through the cut equals the net electrostatic force on the charge

beyond it. The left-hand side is a surface integral of a quantity carrying a *third* derivative of the density; the right-hand side is a volume integral of Coulomb forces. Nothing makes them equal except the identity.

What was measured. On the two-kernel system of Section 4 — unit charges at $z = \pm D/2$, one electron each, split at the midplane, the partner by reflection — with every derivative taken by independent fourth-order differences rather than by the solver’s own operator:

- $Q + V = \lambda$ holds pointwise to 1.4×10^{-3} of the size of Q .
- The integrated balance converges at *fourth* order in the smooth interior — the order of the difference operators — reaching a relative residual of 8.5×10^{-5} .
- At the *free interface* it converges at only *first* order, reaching 1.7×10^{-3} . A control surface cut through the kernel cusp is worse, at about 1%.
- The Hellmann–Feynman force on a kernel and $-dE/dD$ taken from the same code’s energies agree to 5.7×10^{-4} at fourth order, once the finite-difference truncation is removed.

The informative part is the rate, not the pass. That the identity holds was expected — (10) is the chain rule. What was not known in advance is *where* the discretisation of this model loses its accuracy, and the answer is unambiguous: fourth order in the bulk, first order at the free interface. The interface is the accuracy-limiting feature of the whole construction. This matters beyond bookkeeping, because interface motion is exactly where Section 8 looks for dissipation: the dissipative degree of freedom, if candidate 1 is right, lives on the surface which is also the least accurately resolved part of the computation. Any measurement of a viscous stress will therefore have to contend with a first-order error in the same place, and that is worth knowing before the measurement rather than after.

Two cautions about the test itself, which generalise. The kernel softening used throughout this work is $\epsilon = 0.285h$, tied to the mesh. That is a calibration and not a discretisation: it sharpens the kernel as the mesh refines, so the continuum problem moves with h and a mesh ladder measures two things at once. For an identity test ϵ must be held fixed, and was. Second, the energy differences in the force test must be taken on *commensurate* separations $D = (2i+1)h$, which put the kernels on cell centres; off-commensurate differencing produces egg-crate noise the size of the quantity being measured. But commensurability forces the difference width to be at least $2h$, so the narrowest admissible difference still carries an $O(h^2)$ truncation error — which is what the raw 14–36% discrepancy in the force comparison turned out to be, the ratio of the errors at widths $2h$ and $4h$ coming out at 0.19–0.24, i.e. the 1/4 of a squared width. Richardson extrapolation over the two widths removes it and the agreement above is what remains. A force comparison on a fixed mesh is not a test of anything until those two effects are separated.

13 The work programme

In order, with the first two using machinery that already exists. Items 0 and 2 are done.

0. The one-domain identity. *Done* — Section 11.

1. **Two domains with a moving interface.** The free-surface problem in its simplest instance: two densities in contact, interface moving with the normal velocity, chemical potential continuous. Check against the static answer in the limit of vanishing velocity, which is already computed.
2. **The momentum balance as an identity, verified numerically.** *Done* — Section 12. The stress of (7) integrates to the forces the variational solver reports, computed by different code, and the free interface is identified as the accuracy-limiting feature.
3. **The dissipation question,** which after Section 3 is the only part of the problem not already answered in principle. Measure whether interface motion against moving kernels produces a viscous stress; whether its magnitude scales with interface area as candidate 1 predicts; and whether the resulting stress is local and linear in the strain rate, without which it is not Navier–Stokes.
4. **More conservative coefficients, directly.** Extend the sodium calculation of Section 3 to the full elastic tensor of a simple cubic metal, and to a second material, by computing $E(\epsilon)$ on a strained cell. No new machinery; the test is whether the 44% of the sodium bulk modulus is typical or exceptional.
5. **Verify the *form*, not just the coefficient.** Compute $E(\epsilon)$ across a range and check that it is quadratic where linear elasticity is observed and departs from quadratic where nonlinearity is observed. This is the computational substitute for a homogenisation theorem.

14 What would make this a solution, and what would refute it

A programme of this kind is worth only as much as its failure conditions, so they are stated in advance.

It would be *established* if: the one-domain dynamics reproduces free propagation exactly; the two-domain free-boundary problem is well posed and reduces to the known statics; the momentum balance agrees with the variational forces; and a dense-phase equation of state comes out within the model’s usual few per cent.

It would be *refuted*, or at least confined, if: the free-boundary problem is ill posed, which would mean the dynamics cannot be written for more than one electron; or the irrotationality and circulation-quantisation conditions turn out to be incompatible with the non-overlap constraint, which would mean the Madelung form and the multiphase form cannot both hold; or the coarse-grained stress has no dissipative part and finite precision cannot supply one, in which case the programme reaches Euler and stops — a real result, but not Navier–Stokes and not a solution of the problem as posed.

The claim being made at this stage is deliberately modest, and it is a claim about structure rather than about having done the work. The classical route derives continuum mechanics from atomism for a dilute gas, by a limit, after a century of effort. This route starts from a description that is already continuous, so the kinematics and the constitutive law come for nothing, and the entire difficulty is relocated into a single question — where dissipation comes from. That is a better place for the difficulty to be, because it is one question rather than many, and because it can be answered by computation.

References

- [1] Y. Deng, Z. Hani and X. Ma, *Hilbert’s sixth problem: derivation of fluid equations via Boltzmann’s kinetic theory*, arXiv:2503.01800 (2025).
- [2] *Comment on “Hilbert’s Sixth Problem: Derivation of Fluid Equations via Boltzmann’s Kinetic Theory” by Deng, Hani and Ma*, arXiv:2504.06297 (2025).
- [3] C. Johnson, *The Hydrodynamic Form of RealQM: Madelung Hydrostatics, a Surface Tension, and the Missing Half*.
- [4] C. Johnson, *RealQM: Quantum Mechanics as Multiphase 3D Continuum Mechanics*, in [5].
- [5] C. Johnson, *Real Quantum Mechanics* (Body&Soul, Vol. VI), <https://claes542.github.io/RealMolecule/realquantum3.pdf>.
- [6] C. Johnson, *Conduction in RealQM: the Free Boundary as Charge Carrier*, in [5].
- [7] C. Johnson, *Energy without Mass: the Kinetic Coefficient as a Geometric Scale, and Inertia as E/c^2* .
- [8] C. Johnson, *Computational Thermodynamics* (Body&Soul, Vol. V).
- [9] C. Johnson, *The Clay Navier–Stokes Problem and Computational Well-Posedness*.