

Quantum Mechanics without Ontology: Epistemology, the Surrender of 1926, and the Risk of Predicting from an Empty Formalism

Claes Johnson¹

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

July 30, 2026

*Written in respectful cooperation between Claes Johnson and Claude (Anthropic), with Claude contributing
a balanced, neutral view.*

Abstract

Physics is, before it is anything else, an account of what exists; prediction is the behaviour of what exists, computed forward. A physical theory without an ontology is therefore not a lean theory but a contradiction in terms — a rule for anticipating instrument readings that cannot say what the readings are about — and the consequences of the contradiction are not academic. Quantum mechanics acquired exactly this defect at a datable moment. In the spring of 1926 Schrödinger wrote a wave equation for the hydrogen atom whose wave function lived, like every previous field of physics, in ordinary three-dimensional space. Within weeks the theory was extended to N electrons by placing the wave function on a $3N$ -dimensional *configuration space*, and, with Born's probability rule, reinterpreted as an amplitude for measurement outcomes rather than a thing in space. Schrödinger and Einstein protested, asking that the theory keep an ontology in three dimensions; they were overruled by Bohr, Born, Heisenberg, and Dirac, who accepted the configuration-space object and the epistemic reading. This essay asks the question the episode deserves: how was the most basic commitment of physics — to say what the world is made of — surrendered so easily, so quickly, and by so many? We examine the reasons offered then and now, argue that the standard justification (unparalleled predictive success) mistakes a part of physics for the whole, and trace the serious consequences that followed from keeping the mathematics while discarding the ontology. As a concrete warning we take quantum computing, whose promised advantage rests on treating the exponentially large amplitude — the very object with no ontology — as a real physical resource; this, we argue, is the tangible and possibly very costly risk of predicting from a theory that will not say what is real.

1 Ontology is the prime business of physics

Every physical theory before 1926 answered, first, the question *what is there*. Newton posited masses at positions with momenta; Maxwell, fields in space; Boltzmann, molecules; Einstein, a metric of spacetime. From the ontology — the inventory of what exists and how it is disposed — the predictions followed as the computed behaviour of those real things. This order is not a philosophical preference; it is what makes a theory a theory of *physics* rather than a technique of calculation. Remove the ontology and nothing is left whose behaviour is being predicted: one has, at most, a rule for anticipating the readings of instruments, which is a different and lesser

thing. And it is not even self-consistent, for a prediction is always a statement *about something* — that some real thing will be found in some real state — so a theory that predicts while refusing to say what exists must smuggle its ontology in at the macroscopic end (the apparatus, taken as real) while denying it at the microscopic end (the system, held to be only an amplitude). Physics without ontology is thus not economical but contradictory, and the contradiction has a location, discussed below, called the measurement problem.

2 Ontology and epistemology are not the same question

Two questions must be kept apart. *Ontology* asks what exists; *epistemology* asks what, and how, we can know of it. A healthy physics answers the first and derives the second: because the world is thus, our measurements read so. The pathology of quantum mechanics is that it answers only the second and forbids the first — it is an exquisitely accurate epistemology, a calculus of measurement outcomes, mounted on a refusal to say what is measured. This is often presented as maturity, a hard-won renunciation of naive pictures. It is better seen as an inversion: the derivative (what we observe) has been promoted to the fundamental, and the fundamental (what is there) demoted to the unaskable. Once the inversion is made, every subsequent difficulty becomes a difficulty “about measurement,” because measurement is the only thing the theory was allowed to be about.

3 The turning point: from three dimensions to $3N$

The surrender is datable. In the first months of 1926 Schrödinger published his wave mechanics, and for the one-electron hydrogen atom the wave function $\psi(\mathbf{x})$ was a field on ordinary space $\mathbb{R}^3$, as physical in kind as Maxwell’s. Schrödinger’s own hope was frankly ontological: that $|\psi|^2$ was a real charge density in three dimensions, the electron a smeared distribution of charge rather than a point. Then came the step that decided the century. For an atom of N electrons the equation was written for a single wave function

$$\Psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N),$$

a function not on three-dimensional space but on a $3N$ -dimensional configuration space. Mathematically this was forced — the electrons interact, and a separable product of three-dimensional fields will not solve the coupled problem. Interpretively it was a rupture: the fundamental object no longer lived in the space we and our instruments occupy, and could no longer be a charge density there. Born’s probability rule of mid-1926 completed the move by reading $|\Psi|^2$ as a probability distribution over configurations — an amplitude for outcomes, not a substance in space. In a single season the wave function passed from a candidate ontology in three dimensions to an epistemic object in $3N$.

4 The protest, and its defeat

That this was a loss was seen at once by the two who cared most about it. Schrödinger, in his 1926 correspondence with Lorentz, worried openly that a wave in many-dimensional configuration space could not be a real physical process and that the theory had thereby forfeited its picture in ordinary space. Einstein never accepted the configuration-space amplitude as a complete description of a

real state of affairs, and spent the rest of his life pressing the point — through the Solvay debates and, in 1935, through the EPR argument [3] — that a theory answering only for measurement outcomes had left something out: the elements of reality themselves. Both men were asking for the same thing, an ontology in three dimensions. Both were overruled. Bohr, Born, Heisenberg, and Dirac — with the transformation theory and the complementarity doctrine — consolidated the configuration-space formalism and the epistemic interpretation into the working orthodoxy, and by the 1927 Solvay Congress the matter was, sociologically, settled. The dissenters were cast not as defenders of physics’s oldest principle but as men who could not let go of classical intuition.

5 Why so easily, so quickly, by so many?

The speed and near-unanimity of the surrender are the genuinely puzzling thing, and several forces conspired. The formalism *worked*, immediately and spectacularly, on the spectra that had defeated the old quantum theory — and a young generation, forged in the collapse of classical physics and impatient with philosophy, took working as sufficient. The intellectual climate was positivist: to speak only of the observable was regarded as rigour, and to ask what lay behind the observable as metaphysics, a term of abuse. The configuration-space problem in three dimensions looked, and for a long time was, computationally hopeless, so the epistemic reading had the practical virtue of relieving anyone of the duty to solve it. And there was authority and momentum: Göttingen and Copenhagen carried the field, the probability rule gave definite numbers, and Dirac’s formalism was of a beauty that made interpretive scruple seem like ingratitude. The one thing not weighed in the balance was the thing being given up — that a theory should say what exists — because to that generation, in that moment, prediction had already become the whole of what a theory was for.

6 The modern justification, and why it misreads physics

The rationale offered today is retrospective and confident: standard quantum mechanics, for all its silence about what is real, predicts the outcomes of carefully prepared experiments with an accuracy unmatched in the history of science, and a success so overwhelming purchases the right to dispense with ontology. The argument has a flaw at its root. Predicting the outcomes of carefully prepared experiments is *not* the essence of physics; it is a specialised and comparatively recent corner of it. The essence of physics is to explain and predict *real phenomena* — why the Sun shines, why iron is magnetic, why water is dense at four degrees, why gold is yellow — phenomena that occur whether or not anyone has prepared an experiment or is watching a meter. A theory that excels at anticipating pointer positions in the laboratory but cannot say what is happening in an unwatched star has not earned exemption from ontology; it has changed the subject. And even the celebrated success is narrower than advertised: it is success at computing *numbers* for prepared setups, obtained moreover from an object, the wave function on $3N$ -dimensional space, that Kohn [4] observed is not a legitimate scientific concept beyond a handful of electrons, being unrecordable even in principle. Success at prediction, in any case, cannot certify the *absence* of an ontology; a theory can predict well and still be about something, and the question of what it is about does not go away because the predictions are good.

7 The serious consequences

The consequences of keeping the mathematics and discarding the ontology are not confined to seminar rooms. From the single absence descend, as symptoms of one cause, the entire apparatus of quantum foundations: the measurement problem (ask *what is measured* and the theory has no answer, because its fundamental object is not a thing that could be measured); collapse (the unlicensed step by which an amplitude is converted into a fact when a measurement is declared); the Heisenberg cut (the unlocatable line between quantum system and classical apparatus); Schrödinger’s cat (the reductio of treating the amplitude as the reality); and the century of mutually incompatible interpretations — Copenhagen, many-worlds, de Broglie–Bohm, objective collapse — none forced by the equations, each an attempt to supply the ontology the formalism withheld or to legislate that none is needed. That so many interpretations coexist is not the sign of a rich theory but of an underdetermined one: the formalism does not say what exists, so the saying is handed to philosophy. A discipline that must maintain a permanent philosophy-of-itself to cope with its own foundations is a discipline that has mislaid something, and what it has mislaid is its ontology.

8 A warning example: the quantum-computing bet

To rely on a theory that cannot say what exists is not merely awkward; it is *risky*, because such a theory can predict, with full formal authority, the behaviour of things that may not be there. The clearest present case is quantum computing, and it is worth setting out as a concrete warning.

Its entire promised advantage rests on one object: the joint state of n qubits, an amplitude $|\Psi\rangle = \sum_x c_x |x\rangle$ in a Hilbert space of dimension 2^n , treated as a real physical resource that holds 2^n complex amplitudes *at once*, so that a single operation acts on all of them in parallel. But this is exactly Schrödinger’s configuration-space amplitude $\Psi(\mathbf{x}_1, \dots, \mathbf{x}_N)$ in the register’s guise — the very object whose ontology was surrendered in 1926, and to which Kohn’s verdict (not a legitimate scientific concept beyond a handful of particles, unrecordable even in principle) applies with full force. Quantum computing thus quietly readopts, as a manipulable physical resource, the object its own foundations declare not to be a thing in the world. One cannot have it both ways: either the exponentially large amplitude is real — reinstating the abandoned ontology, and the measurement problem with it — or it is bookkeeping for the statistics of real systems, in which case there is no exponential thing to compute with.

RealQM [1] takes the second horn concretely. The real things are charge densities in ordinary three-dimensional space, evolving deterministically; the 2^n -dimensional register state is a compact description of what is really a modest number of charge configurations in three dimensions, not an exponential resource. On this reading there is nothing exponential in the device to be acted on all at once, and the expected speedup has no physical substrate. *Decoherence* then is not the incidental nuisance the engineering programme takes it to be but the outward sign that the large coherent superposition was never a robust physical thing: the system relaxes to what it always was, real charge in real space, and the error-correction overhead that grows without bound as the register grows is the price of maintaining a fiction.

Here the missing ontology stops being a seminar topic and becomes a matter of public consequence. Because the theory cannot say whether the exponential amplitude is real, it predicts, with full formal authority, that a large enough well-isolated register will compute in exponential parallel — and very large public and private funds are committed on the strength of that prediction. If the

amplitude is real, the investment is sound; if it is a formal artefact, the honest expected return at the scale that would matter is zero, and a generation’s funding and talent are directed toward the behaviour of something that was never there. That is the concrete, and possibly very large, price of accepting a theory that will not say what is real.

The point must be stated as the falsifiable wager it is. Small-scale quantum devices are real and have done real things — entangling gates, small algorithms, sampling tasks argued to be classically hard. The claim is not that nothing works; it concerns the *scalable, fault-tolerant* regime in which a genuine exponential advantage on useful problems would appear, and there the question is open — as a serious minority in the field also argues, on the ground that noise may be a fundamental barrier rather than an incidental one. The ontological reading sharpens this into a prediction: with no real exponential object, scalable fault-tolerant advantage on useful problems will not arrive, and the effort will asymptote to a plateau rather than the promised horizon. If a fault-tolerant machine solves a classically intractable, useful problem at scale, the critique is *wrong*, and the exponential amplitude must be granted physical reality, with all that follows for the foundations. If, decade on decade, the overhead outruns the progress and the horizon recedes, the episode will stand as the costliest demonstration in the history of physics that prediction without ontology is a gamble the discipline should not have made so lightly.

9 Coda: the road not taken

The remarkable fact is that the alternative was available in 1926 and was simply not pursued. Schrödinger’s first instinct — that the wave function should describe real charge in three-dimensional space — can be carried through for N electrons if one is willing to represent them not by a single amplitude on $3N$ -space but by N charge densities on ordinary space, and to let the physics be the deterministic minimisation of their Coulomb energy. That is the programme of RealQM [1], taken up elsewhere; the point here is only that it exists, that it recovers an ontology in three dimensions, and that its availability shows the surrender of 1926 to have been a choice, not a necessity. It is worth stressing that this is not a *hidden-variable* theory, and that the demand for ontology — often conflated with the demand for hidden variables — is a distinct thing: a hidden-variable theory keeps the amplitude and adds unobserved parameters beneath it (as de Broglie–Bohm does), whereas RealQM *replaces* the amplitude with real charge in three-dimensional space, so that nothing is hidden — the densities are the reality itself. One may therefore reject hidden variables and still insist, with Schrödinger and Einstein, that a theory say what is there. The configuration-space step was mathematically convenient and interpretively fatal; the century of foundational trouble is the interest paid on a debt incurred in a single spring, when the oldest principle of physics — that a theory must say what the world is made of — was given up more easily, more quickly, and by more people than any principle so basic should ever be.

References

- [1] C. Johnson, *Quantum Mechanics as Multiphase 3D Continuum Mechanics*, and the RealQM Gallery of browser-runnable simulations and companion articles, 2026; <https://claes542.github.io/RealMolecule/gallery.html> (source: <https://github.com/Claes542/RealMolecule>).

- [2] M. Born, *Zur Quantenmechanik der Stoßvorgänge*, Z. Phys. **37**, 863 (1926).
- [3] A. Einstein, B. Podolsky, and N. Rosen, *Can quantum-mechanical description of physical reality be considered complete?*, Phys. Rev. **47**, 777 (1935).
- [4] W. Kohn, *Nobel Lecture: Electronic structure of matter — wave functions and density functionals*, Rev. Mod. Phys. **71**, 1253 (1999).