

The Born–Oppenheimer Rules to Play the Molecular Game

A molecule is a classical roller coaster full of quantum patches. We must only master seven rules to play the game like professional quantum chemists.

by Mario Barbatti 

Mario Barbatti is a theoretical chemist and physicist researching the interactions between light and molecules. He is a professor of chemistry at Aix Marseille University in France and a senior member of the Institut Universitaire de France.

mario.barbatti@univ-amu.fr | www.barbatti.org

Cite this essay as: Barbatti, M. (2026). The Born–Oppenheimer Rules to Play the Molecular Game. Zenodo. <https://doi.org/10.5281/zenodo.19736480>

This work is licensed under CC BY 4.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0>

Born and Oppenheimer, the game masters

The word *molecule* immediately recalls an image of balls representing atoms connected by sticks representing chemical bonds. However, quantum mechanics, the most advanced theory of physics describing radiation and matter at the most fundamental level, [says nothing](#) about such a ball-and-stick molecular structure. Indeed, molecular structures seem not to be an intrinsic property of molecules but rather something that emerges from an approximation proposed by Max Born and J. Robert Oppenheimer in 1927 to deal with the immense complexity of the quantum theory of molecules.

The Born–Oppenheimer approximation is the foundation of molecular science. But it goes beyond to give us molecular structures: it enables a metaphor in which a molecule is a dot moving along a roller coaster track. A chemical reaction happens when this dot effortfully climbs a track hill and slides down to the other side. The Born–Oppenheimer approximation has such a grip on how quantum chemists think about molecules and reactions that, even when it breaks, they struggle to mend it without giving it up. However, has anyone not holding a chemistry degree ever heard of it? The Born–Oppenheimer approximation is

undoubtedly the most influential contemporary scientific theory, entirely unknown to the public.

Max Born (1882-1970) was a German physicist and mathematician recognized as one of the founders of quantum mechanics. In 1926, he formulated the statistical interpretation of the wave function, which is still in use and known as the Born rule. Born was awarded the Nobel Prize in Physics for this fundamental contribution in 1954. Robert J. Oppenheimer (1904-1967) was an American theoretical physicist who, in 1927, was a Ph.D. candidate under Born's supervision in Göttingen. He is primarily famous for his role in the Manhattan Project during World War II, heading the development of the first atomic bombs. Among many contributions, Oppenheimer and his student Hartland Snyder predicted the gravitational collapse that creates black holes.

Still in the early years of quantum mechanics, Born and Oppenheimer [showed](#) that the molecular wave function—the mathematical function describing all properties of a molecule—could be approximated as a simple product of a nuclear wave function and an electronic wave function computed for a fixed nuclear geometry. This approximation stemmed from the enormous mass difference between nuclei and electrons. Each proton and neutron composing a nucleus is 1836 times more massive than an electron. Thus, the swift electrons should only feel nearly frozen nuclei, while the sluggish nuclei should mostly perceive a mean field from all electrons.

Precisely one century after Born and Oppenheimer's work, the separation between nuclear and electronic motions remains the basis of quantum chemistry, the field of chemistry that simulates molecular processes on a computer rather than performing them in the lab. Thanks to decades of advances in computers and algorithms, anyone can go [online](#) nowadays and compute the properties of molecules with dozens of nuclei and a few hundred electrons within minutes.

We can understand how the Born–Oppenheimer approximation works (and when it breaks) without diving into the complex differential equations that underlie it or the computational methods used to solve them. We should regard the Born–Oppenheimer approximation as a roller coaster game. Knowing the seven rules of this game is all we need to get a fair glimpse of how contemporary quantum chemistry describes a molecule.

Laying down the ground rules

A molecule is a stable collection of electrons and nuclei. Electrons are elementary particles with negative electric charges, and nuclei are tight collections of positively charged protons and uncharged neutrons kept together by nuclear forces.

Quantum mechanics describes a molecule via a wave function, a mathematical function that assigns an intensity and a phase to every point in space at every instant of time. The

bigger this intensity, the more likely it is to find an electron or a nucleus at that point. The phase is a property analogous to a color picked from a color wheel, going from red through yellow, green, blue, purple, and back to red. The phase cannot be directly measured, but it is fundamental to correctly describe quantum systems.

The wave function shape and time evolution are determined by the Schrödinger equation, the master equation of quantum mechanics. Nevertheless, solving this equation for a complex system like a molecule is an extraordinary challenge. That is where the Born–Oppenheimer approximation comes in, delivering a practical way to face the problem.

The Born–Oppenheimer approximation splits the molecular Schrödinger equation into two: an electronic and a nuclear Schrödinger equation. In practical terms, a quantum chemist simulates a molecule in two steps. First, she solves the Schrödinger equation for electrons in the presence of fixed-position nuclei. Then, she tackles the nuclear problem using the information from the electrons obtained in the previous step.

Let's focus on the electronic problem for a moment. For instance, if our quantum chemist wants to simulate the hydrogen molecule H_2 , composed of two protons and two electrons, she starts by fixing the position of the two protons. Say, separating them by precisely 1 Å. (Å, the abbreviation for angstrom, is a unit of displacement convenient for molecular sizes. It corresponds to a billionth of a meter or 10^{-10} m.) Then, solving the Schrödinger equation for the electrons, she computes the electronic energy and wave function. And there is no pen and paper. All the work is done on a computer.




Figure 1. Potential energy curve for a molecule with two nuclei.

Next, the quantum chemist changes the separation between the protons to, say, 1.1 Å and repeats the calculations, obtaining a new energy and wave function. Suppose she redoes this procedure, from protons resting almost touching each other to being so far away that their mutual electric repulsion is null. In that case, she will end up with a function $E(R)$ (the expression is read as “ E of R ”), giving the electronic energy E for each proton-proton separation distance R . Such a function is an example of a potential energy curve. This curve

will have a shape like that in Figure 1, with a minimum energy E_{min} corresponding to the equilibrium geometry R_{min} of H_2 .

Now that the quantum chemist knows $E(R)$, she can move to the second part of the molecular problem, describing the nuclei. With the Born–Oppenheimer approximation, nuclear motion simplifies to solving the nuclear Schrödinger equation to obtain a distribution of vibrations and rotations on the potential energy curve $E(R)$. Indeed, it is likely that she will go even further in the approximation and treat the heavy nuclei as point-like classical objects, obeying Newton’s equations of motion.

Without the Born–Oppenheimer approximation, the molecule is no more than a collection of nuclei and electrons described by a single wave function. It forms what the physicist Eric J. Heller aptly portrays as a structureless many-body pea soup (*The Semiclassical Way*, 2018). It is only when we have the potential energy function $E(R)$ that we start to get something that looks like the ball-and-stick models that come to mind when we think of molecules.

In our example, the quantum chemist simulated an isolated H_2 molecule whose nuclear position is entirely described by the relative distance between the two nuclei. To fix the nuclear frame, we must specify three coordinates for an isolated molecule with three nuclei (H_2O , for example). If the molecule has four nuclei (as in ammonia, NH_3), we need 6 coordinates to specify them. The rule goes as $3N-6$ coordinates for a molecule with N nuclei. (Each nucleus needs three coordinates. But if the molecule is isolated and there are no electromagnetic fields around, we can discard three coordinates describing its translation and another three for its global rotation.)

Therefore, the quantum chemist’s job of dealing with large molecules is awfully complicated. Suppose she wants to describe adenine, one of the nucleobases in our DNA. In that case, she must give 39 coordinates to keep adenine’s 15 nuclei fixed. Then, like for H_2 , she should systematically change the nuclear geometry and repeat the electronic calculations. But now, it is not a matter of stretching along a single dimension R to build an $E(R)$ curve. She must create a potential energy *hypersurface* $E(R_1, R_2, ..., R_{39})$ with 39 dimensions! (We simplify the multidimensional notation with a bold-faced $E(\mathbf{R})$.)

Quantum chemists are cleverer than trying this brute-force approach. They have developed many methods to find the most relevant nuclear coordinates describing chemical reactions. Those methods span techniques for connecting different molecular geometries to methods for propagating a swarm of classical trajectories that explore the most accessed configurations. By the end of this technical work, our quantum chemist will likely have only to worry about a few one-dimensional cuts through the multidimensional potential energy surface $E(\mathbf{R})$.

At this point, the molecule became a dot on the graph. A chemical reaction is a pathway bringing this dot from one region of the potential energy surface to another. The probability of it occurring is determined by the height of the potential energy barrier along this path and

by the total energy the molecule can count on (usually determined by temperature). The molecule is now a dot sliding on a roller coaster. That metaphor is so dominant that we no longer consider the ball-and-stick models. We care only about these dots climbing the hills on the potential energy surface $E(\mathbf{R})$. And note the preponderance of geography-borrowed terms such as *pathways*, *surfaces*, and *barriers*.

The pathways on the potential energy surface will help describe how a molecule reacts with another molecule to form a new species, changes into another nuclear configuration, or dissociates, allowing groups of nuclei and electrons to depart. This approach is general; it is equally valid for yeast fermentation, iron rusting in the presence of oxygen, or a protein's conformational change during its biochemical cycle.

With the potential energy surface, we can play the molecular roller *coaster game*. The rules are simple:

1. Each valley corresponds to a different set of molecules.
2. We can only move between two valleys by rearranging the nuclear positions. We cannot add or remove nuclei after the game starts.
3. We can only move between valleys if the $E(\mathbf{R})$ remains smaller than the total energy.

Let's play. An experimental chemist wants to synthesize a new compound. He thinks he can get product **B** if he starts with reactant **A**. However, because the reagents are too expensive, he asks his quantum chemist colleague for help, who computes the potential energy curve shown in Figure 2.




Figure 2. Potential energy surface $E(\mathbf{R})$ connecting a reactant **A** to several possible products **B**, **C**, and **D**. The temperature determines total energy.

The quantum chemist tells him he can get product **B** starting from **A**. Nevertheless, the energy barrier is high, and the chemical reaction will not be efficient. Indeed, she found that the side-product **C** is more likely, as the reaction barrier from **A** to **C** is lower. Now, the experimental chemist knows that if he wants to get **B**, he must add a *catalyst*. This substance

reduces the energy barrier between **A** and **B**. On the bright side, he notices that once **C** is formed, he could use it to get another helpful product, **D**. The quantum chemist bears more bad news. She warns him that the total energy available is insufficient to cross the barrier between **C** and **D**. If her colleague wants to get **D** via this route, he must heat up **C** to raise the total energy bar.

Setting the quantum patches

So far, our nuclei are no more than classical objects closely following Newton's laws of motion. All quantum information is encoded in the electrons through $E(\mathbf{R})$. However, quantum mechanics predicts that the [nuclear wave function spreads](#) into the barriers in the potential energy surface and may even *tunnel* to the other side (note another geography-inspired term). The lighter the nuclei are, the larger the wave function spreads into the non-classical regions. Single-proton nuclei of hydrogen atoms are light enough to exhibit tunneling, shaking the Born–Oppenheimer approximation. Protons prone to tunneling must be factored out of the nuclear wave function and treated quantum-mechanically, like electrons.

Our game gets a new rule to accommodate tunneling:

4. A proton can tunnel to the other valley even if the total energy is slightly lower than the barrier maximum. We know the tunneling probability.

The Born–Oppenheimer approximation has a core assumption: it holds only if the electronic configuration changes smoothly and slowly as the nuclear geometry changes. That is called the *adiabatic approximation*.

The boiling frog fable may help us understand the adiabatic approximation. As the story goes, a frog resting comfortably in tepid water does not notice the water slowly brought to a boil. The unfortunate animal is cooked to death in an *adiabatic process*. Another frog dropped into boiling water will immediately jump out of the pan. The poor animal is hurt, but a *nonadiabatic process* saves it. Like those frogs, molecules can undergo adiabatic and nonadiabatic processes.

Whenever a molecule experiences an abrupt electronic structure change during nuclear distortion (a nonadiabatic change), the Born–Oppenheimer approximation cannot be invoked, and the potential energy surface $E(\mathbf{R})$ is ill-defined.

Did we reach the limit of our metaphor while the roller coaster crumbled into meaninglessness? Not quite. We can still save our roller coaster. But we need to upgrade it.

For each nuclear geometry, a molecule has an associated electronic structure. It may, for instance, have a nonpolar electronic structure in which the electrons are distributed over all nuclei, so that there are no net electric charge differences across the molecule. But it may also

have a *polar* electronic structure, in which the electrons accumulate at one end of the molecule. Suppose the change between these electronic structures is so smooth that the nuclear wave function gradually adapts and slides on a simple pathway on the potential energy surface when the molecule distorts. In that case, the process is adiabatic, and the Born–Oppenheimer approximation is safe.

Nevertheless, there are plenty of situations where the electronic structure changes abruptly, nonadiabatically. There, the Born–Oppenheimer approximation breaks, and we cannot define the function $E(\mathbf{R})$ well. In such cases, the potential energy surface bifurcates; for the same nuclear geometry, we have multiple valid potential energy values. The molecule can slide on different tracks. Which one should it follow? We have to go back to the Schrödinger equation to answer this question.

When the quantum chemist solved the electronic structure, she did it for the lowest possible energy E for the nuclear geometry \mathbf{R} . Nevertheless, the Schrödinger equation has many other valid solutions with higher energies for that same geometry. Then, we should enumerate them. $E_0(\mathbf{R})$ is the lowest energy potential energy surface or the *ground state* (to borrow another geographic term). It is followed by $E_1(\mathbf{R})$, the first excited potential energy surface. Then, $E_2(\mathbf{R})$, $E_3(\mathbf{R})$, and so on. The molecular roller coaster has multiple tracks stacked on top of each other. They are not parallel, though. They may separate, approach, and even touch each other.




Figure 3. Reactant **A** can adiabatically form product **B** or nonadiabatically form product **B***, which is the electronically excited **B**. Alternatively, the reaction can be triggered by shining light on **A** to get **A*** (photoabsorption), which can then form **B*** (adiabatically) or **B** (nonadiabatically). If **B*** is formed, it may spontaneously form **B** by emitting light (fluorescence). In the crossing region, the Born–Oppenheimer approximation is not valid.

When two potential energy surfaces come close together with a slight or null energy separation, the molecule can cross to the other surface by exchanging heat. For instance, a molecule populating the ground state occasionally finds itself in a geometry where the E_0 is near E_1 , like in Figure 3, when molecule **A** in the ground state moves to the right toward the

crossing region. Instead of continuing to slide on $E_0(\mathbf{R})$ *adiabatically* (to form **B**), molecule **A** may jump to $E_1(\mathbf{R})$ via a *nonadiabatic* process, yielding the excited product **B***. Like tunneling, this state crossing is essentially quantum, without any classical analog. We can use quantum mechanics to estimate the probability that a molecule will follow either the adiabatic or nonadiabatic pathway.

Our game just became more complex, albeit more interesting. Here is the supplementary rule:

5. If the track bifurcates during the reaction, the molecule can follow either branch, releasing or absorbing heat with known probabilities.

At room temperature, molecules usually inhabit the ground state—the lowest track. It is there that ordinary chemistry happens. However, excited states can also be populated. Chemical reactions can do this job. Take luciferin dioxetanone found in fireflies, for example. When it breaks into oxyluciferin and CO_2 , the oxyluciferin may be formed in the excited state.

State crossings are not the only way molecules move between states. They can jump from a lower to a higher state by absorbing light. Shine β -carotene with sunlight; it will leap to an excited state by absorbing blue and red wavelengths, letting yellow and orange pass through, giving carrots their typical color. During light absorption, the molecular and the electromagnetic field wave functions entangle. The Born–Oppenheimer approximation breaks again: a simple product of nuclear and electronic functions cannot represent the molecular wave function. Ground and excited states mix, ultimately allowing the jump from the lower to the upper state.

To account for the electromagnetic field, we need a supplementary rule:

6. If the molecule is in a lower track, it can be promoted to an upper track by absorbing light. We know the probability of populating each higher track.

Once in the excited state, a molecule will move on the high-lying track of that state. It will visit new valleys and cross new barriers that do not exist in the ground state. A lot of different chemistry can happen there. DNA, for example, absorbs UV-A wavelengths so efficiently that organisms have multiple mechanisms to minimize the damage of the excited-state reactions. Retinal, a chromophore in our visual cells, undergoes substantial distortions after absorbing light. Such distortions, which are practically impossible in the ground state, trigger the signaling between the eyes and the brain.

But excited states do not remain populated forever. Some molecules move from the excited to the ground state when the tracks approach or touch, turning the excess energy into heat. DNA and retinal are examples where this stabilization by heat release happens within just ten-billionths of a millisecond or one picosecond ($1 \text{ ps} = 10^{-12} \text{ s}$). It is so fast that the DNA chemical composition may have evolved to benefit from this stability against UV damage.

Alternatively, the excited state of a molecule may be depopulated when the electrons *spontaneously* jump down to the ground state, emitting light. For instance, in Figure 3, **B*** can jump to **B** and emit light. We call the spontaneous light emission fluorescence if the electron spin state does not change, and phosphorescence if it does. The excited oxyluciferin in the fireflies returns to the ground state via fluorescence, creating their extraordinary luminescence.

Spontaneous light emission is another purely quantum phenomenon. It is caused by random fluctuations in energy at each point in space, mixing the molecular and electromagnetic field wave functions. And once more, the Born–Oppenheimer approximation breaks.

Such quantum fluctuations add another rule to the game:

7. If the molecule is in an upper track, it can spontaneously go to the lower track by emitting light. We know the probability of this emission happening.

The wavelength (or color) of the light a molecule absorbs or emits is precisely determined by the energy difference between the two states, spanning the infrared, visible, ultraviolet, and even X-rays. Each molecule has a specific spectral signature (the amount of light it emits or absorbs at each wavelength), which can be used to identify unknown compounds, for instance, in the atmospheres of exoplanets. Experimental techniques have advanced to the point that it is possible to monitor which wavelengths are absorbed or emitted during chemical reactions occurring within 0.1 ps.

We have the rules of our roller coaster game laid out. In a popular version of this game, formulated by John C. Tully and called surface hopping, we use Newton’s equations to predict the nuclear geometry as a function of time as the molecule slides on a single potential energy surface. But at each step, we roll the dice to see whether the molecule tunnels to another valley or jumps to another surface via light absorption, light emission, or heat exchange. Playing this game with advanced computer programs has given quantum chemistry an essential role, helping explain chemical processes in the lab and nature.

Beyond Born and Oppenheimer

Great scientific metaphors often accompany mature science. Such is the case with the roller coaster of quantum chemistry (sketching molecular reactions), the rubber sheet in general relativity (illustrating how a mass distorts spacetime), or the tree of life in biology (drawing the ancestral relationships among all organisms). All these visual metaphors project abstract concepts (energy, multidimensional realities, speciation) onto the conventional space, enabling the scientist to mentally manipulate highly abstract scientific concepts as if they were concrete objects. The scientific metaphors go well beyond didactic roles; they are part of the scientific routine, as when we inquire whether a barrier in the potential energy would

prevent a chemical reaction. They push the boundaries of discovery, such as when we ask what happens to the potential energy surface when the adiabatic approximation is no longer valid.

Naturally, as Thomas Kuhn realized in *The Structure of Scientific Revolutions* (1962), there is always the risk that the scientific community gets so attached to a paradigm that it becomes exceedingly conservative. It will favor patching the metaphor with new features (like the bifurcating stacked tracks in the potential energy roller coaster) over opening space for a novel paradigm that paints new metaphors. But is there any alternative to the Born–Oppenheimer paradigm?

In 1975, Geoffrey Hunter introduced a completely distinct approach to separate the electronic and nuclear wave functions, which avoids the Born–Oppenheimer approximation altogether. This approach began to gain momentum in 2010 when Ali Abedi, Neepa T. Maitra, and Eberhard Gross formulated a time-dependent version, which they named [Exact Factorization](#). In Exact Factorization, the molecular wave function still looks like a simple product of a nuclear and an electronic wave function. But its similarity to the adiabatic Born–Oppenheimer approximation is deceiving. The Schrödinger equation governing the wave function’s time evolution in the Exact Factorization contains terms that ensure its validity in nonadiabatic regimes.

Exact factorization ruthlessly demolishes the Born–Oppenheimer roller coaster. It replaces it with a time-evolving potential energy surface. There are no excited states, nor are there any corresponding stacking tracks. Instead, it offers a dynamic roller coaster that changes shape, altering its minima and barriers during the chemical reaction. It is still too early to assess whether Exact Factorization will become a standard tool for quantum chemists.

Quantum chemistry thrives with new method developments, such as accelerating discovery via artificial intelligence and assembling new algorithms for the coming quantum computing revolution. Except for a few exceptions (such as Exact Factorization), most of these developments still rely on the Born–Oppenheimer approximation, with nonadiabatic patches when needed. Therefore, it is inevitable that its powerful but ludic roller-coaster game metaphor will continue to guide quantum chemists in the foreseeable future.