Toward an Evolutionary History of Terrestrial Oxygen

Impact Mixing, Reservoir Formation, and the Limits of Historical Reconstruction

Flyxion
Independent Researcher

October 2026

Abstract

A recent high-pressure experimental study reports that silicate melts can hold several percent by mass of oxygen beyond what cation valence balance allows, and proposes that ascending magma may release this oxygen and oxidize the shallow Earth. The result invites an obvious question about where such oxygen-rich reservoirs came from, and one candidate answer runs through the Moon-forming impact and any water that Earth possessed before it. This essay uses that question as a worked case for a more general problem in the reconstruction of unobserved planetary histories. Hypotheses about a single contingent event are routinely ranked by how wide a range of initial conditions they accommodate, and the wider hypothesis brings a particular kind of relief, because the event no longer appears to require a coincidence. Breadth of accommodation is nonetheless a property of a hypothesis, whereas what occurred is a property of the world, and the two come apart. The essay sets out the structure of that distinction, applies it to the family of Moon-formation models and to the excess-oxygen result, and lists the observations that would separate the competing accounts of how Earth's oxidized reservoirs arose. The pre-Moon ocean proposal is stated as one entry in a ledger of rival hypotheses. No evidence identified here specifically favours that sequence over competing reservoir histories, and the essay says so throughout.

1 The relief of a wide hypothesis

Some explanations satisfy before they have been tested. A phenomenon that looked like a knife-edge coincidence is shown to follow from a broad family of circumstances, and the sense that something improbable has been made ordinary arrives at once. Nietzsche described this relief in connection with the reduction of the unfamiliar to the familiar, which he took to be a source of comfort and a feeling of power well before it is a source of knowledge [10, 11]. The present essay applies that observation to a case in which the relief is generated by a legitimate scientific virtue. When one account of an event requires fine tuning of its initial conditions and a rival account tolerates a wide range, the rival often earns more support, although a guaranteed ordering exists only when the admissible set of the wide account contains that of the narrow one, and in general the comparison depends on the prior each model assigns to its own parameters and on the likelihood of the observations given those parameters (Appendix A). Support under a stipulated prior is not the identification of the history that occurred. The

satisfaction produced by breadth tends to be reported as though it were evidence about what happened, and the move from the one to the other is the subject of what follows.

The motivating example is the origin of the Moon. The long-standing difficulty with giant-impact accounts has been that the angular momentum, composition, and mass of the Earth–Moon system seemed to require a narrow set of impact conditions, whereas models in which a high-energy collision produces a vaporized, rapidly rotating, doughnut-shaped structure widen that set considerably [9]. Preference for the wider model is reasonable. It would be a mistake, however, to say that the wider model has thereby been shown to describe the event, since the event was a single collision with a particular geometry, a particular number of impactors, and a particular energy, and none of these quantities has been measured. What the wider model establishes is that a Moon-forming outcome is robust across many conceivable histories, which is a statement about the space of histories and not about the one that was realized.

A second line of work supplies a different kind of example. A study published in *National Science Review* reports that fully molten basaltic glass quenched from high-pressure experiments contains oxygen in excess of the amount balanced by the measured cations, reaching about 5.2 ± 0.2 wt.% at 7 GPa under strongly oxidizing conditions [16]. The authors draw out a possible link between oxygen-rich deep melts and the oxygenation of the atmosphere. If one grants the experimental result its full weight, a further question follows about how an oxygen-rich region of the mantle could have formed, and the answers on offer are numerous. Among the candidates is the possibility that water present on Earth before the Moon-forming impact was vaporized, mixed with silicate material, and returned to the planet in an altered oxidation state. That proposal is the kind of hypothesis the essay is concerned with: it is coherent, it connects two bodies of work, it cannot at present be distinguished from several rivals, and it is attractive partly because it makes a large number of events fall into place.

2 What the excess-oxygen result establishes and what it leaves open

The experimental claim is carefully bounded in the source. The authors define excess oxygen as the directly measured oxygen content of a glass minus the oxygen implied by assigning stoichiometric oxide proportions to all measured cations, with water, carbon dioxide, and chlorine corrections applied [16]. They assume that all iron is trivalent, which is the most conservative choice because any divalent iron would raise the computed excess. Under that convention the excess in fully molten basalt rises from approximately zero at one atmosphere to 5.2 ± 0.2 wt.% at 7 GPa when platinum dioxide supplies oxygen, while comparable runs without the added oxygen source give values between 0.2 ± 0.7 and 0.5 ± 0.4 wt.%. Molecular dynamics simulations of a magnesium silicate composition reproduce the sign and general trend of the pressure dependence, giving 0.33 ± 0.17 wt.% at one atmosphere and 2.12 ± 0.26 wt.% at 9 GPa, with a small population of oxygen atoms whose computed charges resemble those of peroxide species. Raman spectra of the quenched glasses show no molecular oxygen band, which the authors read as structural incorporation.

The authors are explicit that the platinum dioxide imposes an end-member condition and that the highest values represent a capacity, not an expected abundance in mantle-derived magma. They do not claim that mantle magmas generally contain significant excess oxygen. Their illustrative calculation, in which release of roughly 0.27 wt.% excess oxygen from the magma volume of large igneous provinces emplaced between about 2.5 and 2.0 billion years ago would match the oxygen inventory of the present atmosphere, is described as a scaling exercise. The atmospheric comparison can be reproduced from the stated inputs only if the quoted volume is converted from an extrusive volume to a total magma volume. The quoted 7.2 million cubic kilometres corresponds to 41 provinces of 0.18 million square kilometres at 1 km thickness, and the source gives an extrusive-to-intrusive ratio of 5%. If that ratio is applied, the total is between about 144 million cubic kilometres (extrusive as 5% of the total) and about 151 million (extrusive as 5% of the intrusive volume), depending on the convention. Either gives a mass near 4.3 to 4.6×10^{20} kg at the stated density, of which 0.27 wt.% is about 1.2×10^{18} kg, close to the mass of molecular oxygen in the modern atmosphere. If the quoted volume were instead taken as the total, the same fraction would supply only about one twentieth of that mass. The text does not spell out the volume convention, and the result depends on it by a factor of twenty (Appendix B). The defensible observation is that reproducing the comparison requires a convention that the source leaves insufficiently explicit.

Three features of the result matter for the historical question. First, nothing in it demonstrates excess oxygen in any natural sample, so the existence of oxygen-rich mantle reservoirs is a premise supplied from earlier literature on mantle redox heterogeneity and not an output of the experiments. Second, the quantity measured is a systematic residual obtained by subtracting a large calculated oxygen total from a large measured one, so calibration offsets in the analysis of oxygen, underestimated hydrogen, or unmodelled valence states of other elements would each appear as apparent excess. The main text addresses part of this by reporting no correlation of excess with measured water and carbon dioxide, by running oxygen-source-free controls, and by the conservative iron assumption, but it does not quantify the matrix effects that arise when oxygen is measured by electron microprobe in glasses quenched from very different pressures, and the simulated one-atmosphere value does not coincide with the experimental one. Third, the authors note that the values are minima because oxygen may have degassed during quenching, which cuts in the opposite direction. A responsible reading holds both facts at once: the effect is probably real in the sense of being larger than the stated uncertainty in oxygen-rich runs, and its magnitude and its natural relevance are open.

3 Oxygen atoms and oxidizing power

Any account of the origin of an oxidized reservoir has to be stated in terms of oxidizing power, not oxygen atoms. Oxygen is the most abundant element in the silicate Earth, and one modern ocean alone contains roughly 1.2×10^{21} kg of oxygen bound in water, about a thousand times the oxygen in the present atmosphere. Water brings its oxygen together with hydrogen, however, and the pair carries no net oxidizing power unless the hydrogen is subsequently lost or sequestered in a reduced reservoir. For each mole of molecular hydrogen that leaves the

system, two moles of electron-equivalents of reducing power go with it, so that the remaining material is more oxidized by the amount that half a mole of molecular oxygen would supply. The decisive variable in any pre-Moon ocean hypothesis is therefore the fate of the hydrogen, and the relevant literature on steam atmospheres and hydrogen escape after large impacts is the place where the hypothesis is most exposed [6, 15].

A second constraint concerns the channel through which oxidizing power is stored. The conventional account places it in the ferric-to-ferrous ratio of the mantle [5], and the excess-oxygen paper explicitly excludes that channel by computing its excess after assuming all iron is already trivalent. Computing the excess with all iron assigned the +3 state shows that the measured oxygen exceeds that particular stoichiometric accounting, and it does not show that a natural reservoir must exhaust its ferrous iron before non-stoichiometric oxygen can occur. Whether the two channels coexist, compete, or exchange depends on equilibrium, oxygen fugacity, melt composition, and kinetics. For the pre-Moon ocean proposal this becomes an open question for a coupled chemical model: how much ocean-derived oxidizing power would be taken up by iron, how much would be retained in melts as excess oxygen, and how much would be exported by other channels, none of which can be answered from the budget available here.

4 The Moon-forming impact as a contingent event

The canonical giant-impact model supposes a Mars-sized body striking the proto-Earth obliquely and at modest speed [2]. Subsequent work widened the family. A small impactor striking a rapidly rotating Earth can produce the observed system if later dynamical evolution removes angular momentum [4], an impact between bodies of similar mass can mix the silicate reservoirs more thoroughly [3], and a sequence of smaller impacts forming several moonlets can assemble a single satellite [14]. The synestia model describes the outcome of high-energy, high-angular-momentum collisions as a single hot, vapor-rich structure in which the planet and the disk are not separate objects, and shows that a Moon can condense within it [9]. The model is preferred partly because it accommodates a much wider range of impact conditions than the canonical scenario, and partly because the mixing and equilibration it implies is consistent with the near identity of Earth and lunar isotopic compositions [12, 17]. These are two distinct reasons, and the first is the one that produces relief.

The record that can discriminate among these histories is thin. The lunar cratering record, including the interval of heavy bombardment near 3.9 billion years ago whose character is still debated, preserves information about later impactor populations and has little to say about the formation event itself, which occurred several hundred million years earlier. The constraints on the event are the total angular momentum of the Earth–Moon system, the mass and bulk composition of the Moon, the small iron core, and the isotopic similarity of the two bodies. Each of these is compatible with several impact geometries, so that the impact angle, the impactor mass, and even the number of impactors remain undetermined. It is possible that one body struck once, that two bodies struck in close succession, or that several smaller bodies struck in sequence, and the observed system does not currently discriminate among these. These scenarios are not interchangeable by dividing one impact's energy among several,

since timing, angular momentum, material delivery, and satellite dynamics differ between them, and each must be modelled separately.

Under such conditions the procedure of inquiry is pruning. Hypotheses that require improbable fine tuning are retained with low weight, and hypotheses that accommodate many histories are retained with higher weight. This is sound practice, and it is also incomplete as a description of what has been learned. Under a stipulated prior the weights are posterior probabilities that the realized conditions lay in the set a hypothesis admits, and they are only as good as that prior. What they do not provide is certainty, a verdict independent of the prior, or the identification of a unique history. The truth of each hypothesis was fixed by the event, and observation changes what is known about it and not what it was. Appendix A states this formally. A hypothesis with substantial posterior support can be false, because the realized history might have lain in a region it excludes, and a hypothesis with little support can be true.

5 Hypotheses that are not built

A further selection operates before pruning begins. Consider the possibility that Earth possessed an ocean before the Moon-forming impact, that the impact vaporized it, and that the vapor was incorporated into the planet and the disk. One can extend the scenario to two or three earlier oceans, each removed by an earlier collision. Nothing in the available record excludes the extension, and nothing in the available record distinguishes it from the single-ocean case or from the case of no ocean at all. Such hypotheses are rarely constructed, and the reason is not that they are judged false. A hypothesis that accommodates every geometry and every number of impacts forbids almost nothing, and a hypothesis that forbids almost nothing cannot be tested, so effort is directed toward hypotheses that make differential predictions [13]. The preference expressed by that practice follows testability and tractability, and it should not be mistaken for a finding about the history of the planet.

This produces a tension that the essay accepts and does not attempt to resolve. Breadth of accommodation raises the weight of a hypothesis, and excessive breadth lowers its informativeness. A wide synestia hypothesis is preferred over a narrow impact hypothesis, and an unrestricted hypothesis about repeated ocean loss is declined, and the difference between the two is the degree to which each still forbids some outcomes. The relief that attends a wide hypothesis is evidence that it is easy to hold, and the discipline of the practice consists in recording that the ease is not an observation.

6 A ledger of hypotheses for oxidized reservoirs

The excess-oxygen result needs a source of oxidized material in the mantle, and the literature provides several candidates. Table 1 sets out the ones relevant here in a common format so that their commitments can be compared directly. The pre-Moon ocean entry is included on the same footing as the others, and no claim is made that it ranks above them.

The hypotheses are not exclusive, and a plausible history may combine several of them, which is itself a reason for caution in assigning explanatory credit to any one. The pre-Moon

ocean entry carries a particularly long chain of dependencies. It requires that an ocean existed, that its water was vaporized and mixed with silicate material, that hydrogen subsequently left, and that the oxidized residue reached depths where melts can hold excess oxygen, which are pressures of several gigapascals and far above the tens of bar reported for the lunar-forming region of the synestia. Each link is individually conceivable, and the conjunction is weaker than any single link. The proposal is also complicated by the question of whether an intact earlier ocean is the right object at all, since a collision of the required energy would leave only a volatile inventory and an oxidation state, so that the hypothesis is better stated as inheritance of volatiles than as survival of an ocean. The remaining hypotheses also carry costs. Radiolysis is a well-documented process in crustal settings [8], but its extrapolation to a deep reservoir requires a quantitative budget that has not been given, and the explanation of reservoir formation through mineral evolution is best restricted to deep-mantle mineralogy, since much of the surface mineral record is more plausibly a consequence of oxygenation than a cause of it.

7 What could discriminate among the hypotheses

Several classes of observation could distinguish the entries in the ledger, and a hypothesis that predicts the same observables as another should be marked as not yet distinguishable. The first class is the direct measurement of excess oxygen in natural samples, which would test H0a, H0b, and H1 and would fix the magnitude of the effect in melts derived from the mantle. The second is hydrogen and oxygen isotope composition of deep-sourced materials, which bear on the origin of water and on whether hydrogen was lost, a point on which H5, H6, and H7 diverge. The third is noble gas signatures of ancient mantle reservoirs, which constrain whether any primitive domain survived to the present and so bear on H2. The fourth is the redox record of mantle-derived rocks of different ages, which tests whether oxidized domains were present when the large igneous provinces were emplaced. A fifth is the estimation of the budget itself, including the oxidizing power needed to produce an excess of the stated size in a given volume of mantle and the oxidizing power available from each source (Appendix B). Where several hypotheses give the same prediction for all of these classes, the correct statement is that the history is underdetermined, and that wider hypotheses remain preferred for reasons of probability and not of evidence.

8 Provisional history as a form of knowledge

A history of terrestrial oxygen assembled from this material would have the character of a constrained reconstruction in which the mechanisms are physically admissible, the constraints exclude some sequences, and the particular sequence remains undetermined. That is a respectable and useful form of knowledge, and it differs from the narrative form that explanatory relief tends to encourage. A narrative runs from an early ocean, through an impact, to an oxidized mantle and an oxygenated atmosphere, and each step is made plausible by the preceding one. A reconstruction in the sense intended here retains the branching, assigns each branch

the evidence that bears on it, and declines to collapse the tree where the evidence does not warrant it. The wide hypothesis then has its proper role as a guide to where search should be concentrated, and the single realized history stays a matter for observation.

The excess-oxygen study is a good test of this discipline precisely because its result is attractive. It offers a mechanism linking deep magmatism and surface oxidation, and the time coincidence between large igneous provinces and oxygenation episodes lends the link additional appeal. The appeal is not evidence. A reader who finds the account satisfying has learned something about the account, and the proper response is to ask what observations would be different if it were false. The essay's own proposal about a pre-Moon ocean is subject to the same test. Individual mechanisms in its chain have independent support, such as vaporization and mixing in high-energy impacts and the loss of hydrogen from steam atmospheres, but no evidence identified here specifically favours their conjunction over competing reservoir histories, and no observation conflicts with it. That is the condition in which hypotheses are formed and not the condition in which they are accepted.

A Breadth, support, and the realized history

Let Θ be the space of possible initial conditions of the event (impact angle, velocity, impactor masses, number of impactors, pre-impact state of the planet), and let O denote the observations. A hypothesis H is the proposition that the realized conditions θ^* lie in an admissible set $A_H \subseteq \Theta$. Under a probability measure μ on Θ expressing a prior over histories, and writing $C_O \subseteq \Theta$ for the conditions that reproduce O, with $\mu(C_O) > 0$, the posterior probability of the proposition is

$$w(H \mid O) \;=\; \frac{\mu(A_H \cap C_O)}{\mu(C_O)}.$$

This is a genuine conditional probability, and it expresses support under the stipulated measure. It does not provide certainty, a verdict independent of μ, or the identification of a unique history. The truth of H is the indicator $\mathbf{1}[\theta^* \in A_H]$, which was fixed by the event; observation changes what is known about it and not its value. A hypothesis with w near one can be false, and a hypothesis with w near zero can be true.

Set inclusion gives a guaranteed ordering. If $A_H \subseteq A_{H'}$, then $\mu(A_H \cap C_O) \leq \mu(A_{H'} \cap C_O)$ under every probability measure, so the wider hypothesis cannot receive less support. For competing physical models that are not nested, an informal comparison of breadth gives no such guarantee, because each model ordinarily carries its own parameter distribution π_H over Θ_H and its own likelihood, and its support is governed by the marginal likelihood

$$P(O \mid H) \;=\; \int_{\Theta_H} P(O \mid \theta, H)\, \pi_H(\mathrm{d}\theta).$$

A model that permits many initial conditions gains support when a substantial share of its prior probability leads to the observed outcome. Merely permitting more conditions does not guarantee that gain, because probability spread over conditions that do not reproduce O dilutes the marginal likelihood. This is the precise sense in which the relief of a wide hypothesis outruns

the support it has earned.

Informativeness runs in the opposite direction. Writing F_H for the set of observation outcomes incompatible with every $\theta \in A_H$, the inclusion $A_H \subseteq A_{H'}$ implies $F_{H'} \subseteq F_H$. A hypothesis with $A_H = \Theta$ has $F_H = \varnothing$ and is untestable. Hypotheses about multiple earlier oceans approximate this limit with respect to available observables, which accounts for their omission.

B An oxidizing-power budget

Define oxidizing power in electron-equivalents, with one mole of molecular oxygen counting as four. Loss of one mole of molecular hydrogen from a closed system containing water leaves the remainder more oxidized by 2 electron-equivalents, so for $n_{\mathrm{H_2}}$ moles lost,

$$\Delta\Omega_{\mathrm{residue}} \;=\; 2\, n_{\mathrm{H_2}}.$$

Oxidizing Fe^{2+} to Fe^{3+} gives up one electron per iron atom and therefore consumes one electron-equivalent of supplied oxidizing power. Taking x_{Fe} to mean the mass fraction of the mantle that is initially iron present as Fe^{2+}, with molar mass m_{Fe}, the oxidizing power the iron channel can consume in a mantle of mass M is at most

$$\Omega_{\mathrm{Fe}}^{\mathrm{max}} \;=\; \frac{M\, x_{\mathrm{Fe}}}{m_{\mathrm{Fe}}}.$$

The supplied oxidizing power is partitioned among all channels that can take it up,

$$\Omega_{\mathrm{supplied}} \;=\; \Omega_{\mathrm{Fe}} + \Omega_{\mathrm{EO}} + \Omega_{\mathrm{other}} + \Omega_{\mathrm{export}},$$

where $\Omega_{\mathrm{Fe}} \le \Omega_{\mathrm{Fe}}^{\mathrm{max}}$, Ω_{EO} is the portion retained as excess oxygen in melts, Ω_{other} covers further sinks such as other multivalent elements, carbon, and sulfur, and Ω_{export} covers loss from the system. The partition is not sequential. How the supply divides among the channels is determined by equilibrium, oxygen fugacity, composition, and kinetics, and it is the quantity a coupled chemical model would have to supply.

For a melt of mass m_{melt} at excess-oxygen mass fraction ε, the quantity $4\, m_{\mathrm{melt}}\, \varepsilon / m_{\mathrm{O_2}}$ is an O_2-equivalent potential, meaning the oxidizing power the excess would deliver if it were released as molecular oxygen. Its actual electron-accepting capacity depends on the chemical state of the oxygen and on the reference products, which the experiments do not fix. Each hypothesis in Table 1 supplies a different expression for $\Omega_{\mathrm{supplied}}$, and comparing those expressions with the potential needed for a given release is the quantitative test of whether the hypothesis can matter.

As a check on the scaling quoted from the source, the atmospheric oxygen mass is approximately 1.2×10^{18} kg. Applying the 5% extrusive-to-intrusive ratio to the 7.2 million cubic kilometre extrusive volume gives a total of about 144×10^6 km^3 if the extrusive volume is 5% of the total, or about 151×10^6 km^3 if it is 5% of the intrusive volume. At 3010 kg m^{-3} these are about 4.3×10^{20} and 4.6×10^{20} kg, and 0.0027 of those masses is about 1.2×10^{18} kg in

either case. If the quoted volume were instead the total, the product would be about 5.9×10^{16} kg. The comparison therefore holds under either extrusive-to-total convention and fails if no conversion is made, and the source does not state which convention is intended.

References

[1] Andrault, D., Muñoz, M., Pesce, G., et al. (2018). Large oxygen excess in the primitive mantle could be the source of the Great Oxygenation Event. *Geochemical Perspectives Letters*, 6, 5–10.

[2] Canup, R. M., and Asphaug, E. (2001). Origin of the Moon in a giant impact near the end of the Earth's formation. *Nature*, 412, 708–712.

[3] Canup, R. M. (2012). Forming a Moon with an Earth-like composition via a giant impact. *Science*, 338, 1052–1055.

[4] Ćuk, M., and Stewart, S. T. (2012). Making the Moon from a fast-spinning Earth: a giant impact followed by resonant despinning. *Science*, 338, 1047–1052.

[5] Frost, D. J., Mann, U., Asahara, Y., and Rubie, D. C. (2008). The redox state of the mantle during and just after core formation. *Philosophical Transactions of the Royal Society A*, 366, 4315–4337.

[6] Hamano, K., Abe, Y., and Genda, H. (2013). Emergence of two types of terrestrial planet on solidification of magma ocean. *Nature*, 497, 607–610.

[7] Hu, Q., Kim, D. Y., Yang, W., et al. (2016). FeO_2 and FeOOH under deep lower-mantle conditions and Earth's oxygen–hydrogen cycles. *Nature*, 534, 241–244.

[8] Lin, L.-H., Slater, G. F., Sherwood Lollar, B., Lacrampe-Couloume, G., and Onstott, T. C. (2005). The yield and isotopic composition of radiolytic H_2, a potential energy source for the deep subsurface biosphere. *Geochimica et Cosmochimica Acta*, 69, 893–903.

[9] Lock, S. J., Stewart, S. T., Petaev, M. I., et al. (2018). The origin of the Moon within a terrestrial synestia. *Journal of Geophysical Research: Planets*, 123, 910–951.

[10] Nietzsche, F. *The Gay Science*, §355.

[11] Nietzsche, F. *Twilight of the Idols*, "The Four Great Errors", §5.

[12] Pahlevan, K., and Stevenson, D. J. (2007). Equilibration in the aftermath of the lunar-forming giant impact. *Earth and Planetary Science Letters*, 262, 438–449.

[13] Popper, K. R. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.

[14] Rufu, R., Aharonson, O., and Perets, H. B. (2017). A multiple-impact origin for the Moon. *Nature Geoscience*, 10, 89–94.

[15] Sossi, P. A., Burnham, A. D., Badro, J., et al. (2020). Redox state of Earth's magma ocean and its Venus-like early atmosphere. *Science Advances*, 6, eabd1387.

[16] Xu, Y., Niu, H., van Westrenen, W., Liu, R., Wu, P., Mao, H.-K., and Lin, Y. (2026). Magma as a carrier of non-stoichiometric oxygen linking Earth's interior and atmosphere. *National Science Review*, 13, nwag539.

[17] Young, E. D., Kohl, I. E., Warren, P. H., et al. (2016). Oxygen isotopic evidence for vigorous mixing during the Moon-forming giant impact. *Science*, 351, 493–496.

Label	Source of oxidizing power	Required steps beyond the experiment	Potential discriminants and limiting observations
H0a	None; the excess is an analytical or experimental artifact	None	Independent oxygen analysis and controlled fugacity that reproduce the residual would count against it, and one that removes it would support it
H0b	None; the excess is real but confined to strongly oxidizing laboratory conditions	No natural melt reaches the required oxygen fugacity and pressure	Excess found in natural glasses or melt inclusions, with iron valence measured directly, would count against it
H1	Ordinary Fe^{2+}/Fe^{3+} redox, with no independent oxygen component	Mantle redox variation follows iron alone	Natural samples showing oxygen beyond what measured iron valence accounts for
H2	Primitive mantle oxygen excess, retained from accretion [1]	Survival of unhomogenized domains through mantle convection	Mantle domains of the required age show uniformly reduced signatures
H3	Recycled surface-oxidized material delivered by subduction	Transport of oxidized material to melting regions	Oxygen-rich plume sources lacking a recycled isotopic component, which counts strongly only if the hypothesis predicts that such a signature is preserved and detectable
H4	Iron disproportionation with removal of metal to the core [5]	Early metal segregation sufficient to oxidize the residual mantle	Core composition or mantle redox inconsistent with the required metal fraction
H5	Water and iron reacting at depth to form iron peroxides with hydrogen loss [7]	Stability of the phases and escape of hydrogen from the lower mantle	Absence of the phases or of hydrogen-depleted signatures in deep-sourced material
H6	Water inherited from a pre-Moon ocean, processed through the giant impact	An earlier ocean, hydrogen loss after the impact, return of oxidized material to the deep mantle at pressures far above those of the disk	Absence of an ocean-compatible volatile signature in early mantle reservoirs, or evidence that hydrogen was retained, bears against it to the extent that the hypothesis predicts such signatures survive
H7	Radiolysis of water with separation of the products [8]	Sustained separation against recombination and reaction with reduced minerals	Insufficient radiogenic energy over the relevant time for the required oxidant mass

Table 1: Competing sources of oxidizing power for oxygen-rich mantle reservoirs. The entries are not mutually exclusive, and several could contribute.