

CENTAUR RESEARCH

A Bayesian Analysis of Life After Death

Conscious Continuation, Brain Dependence, and Alternative Architectures

Rodney Malpert, J.D., M.A. (Cornell University)

RodMalpert.com

With ChatGPT (OpenAI; GPT-5.6 Sol and GPT-6.1 Sol). Earlier adversarial contributions from Claude (Anthropic) and Grok (xAI).

October 6, 2026

Abstract

Does irreversible biological death end conscious existence, or can a subject or a meaningful continuation of its experiential history persist? This paper examines both empirical survival claims and competing causal architectures. It separates dependence on the brain from the further claim that the brain produces the whole subject, and it distinguishes a generated simulated mind from an external subject experiencing a simulated life.

The model compares four local architectures: locally produced subjects with no specified continuation mechanism; participatory simulation; nonsimulation idealist interfaces; and nonsimulation separable or transmission theories. It updates their starting weights on brain dependence, assigns a continuation conditional to each, and applies a modest, explicitly discounted empirical update. These are elicited judgments, not measured odds or an academically agreed probability.

With the declared low-preservation inputs, structural continuation is about 44% and the empirical benchmark about 47%. The expanded generated-mind model in Appendix D gives 48.8% at a 25% generated-mind continuation conditional and 51.2% at 50%. These are scenarios, not measured rates; no new preferred conditional is inferred from the interview. The broad 20-70% sensitivity envelope is not a confidence interval.

The direct paranormal record does not drive the estimate, although selected corroborated NDE perceptions receive a modest positive allowance. Near-death experiences occur in people who survived resuscitation, rather than after established irreversible death. Prospective target tests have limited power; stronger replication studies weaken specific precognition claims; and childhood past-life reports require case-specific leakage and documentation audits. Participatory simulation is a coherent possibility, but its prevalence, observer measure, motives, and continuity conditions are unmeasured. The paper's central contribution is to expose these assumptions rather than turn their logical possibility into proof of survival.

Executive Summary

- The low-preservation benchmark gives about 47% broad continuation. Appendix D makes generated-mind preservation explicit; its sensitivity crosses 50% without changing other inputs.
- Detailed brain dependence remains the strongest conventional evidence. Anesthesia is part of that cluster, rather than an additional independent penalty.
- Participatory simulation connects a pre-existing subject to an avatar. Generated simulation creates a subject within the implementation. Neither architecture guarantees continuation without further assumptions about preservation and identity.
- Selected corroborated NDE perceptions receive a separate modest positive allowance. Documentation quality and ordinary information access remain central limitations; the factors are judgments, not estimated effect sizes.
- Strong computational functionalism has a working weight of 25%, with substantial uncertainty. Rejecting that strong thesis does not establish that all artificial systems are unconscious.
- Idealism and dualism must be divided into families with different continuation implications. Strict epiphenomenal property dualism can terminate experience with brain function.
- Starting architecture weights and participatory frequency matter much more than the small empirical update. Low participatory frequency alone reduces the worked estimate to about 34%; a lower idealist and dualist starting allocation reduces it to about 36%.
- These numbers describe this model. They do not establish a lower bound that every rational observer must assign to survival.

1. The Question: What Counts as Life After Death?

The paper does not ask whether every detail of a human autobiography survives unchanged, nor whether an immaterial Cartesian soul literally detaches from the corpse. Those are particular mechanisms. The target proposition is broader: biological death fails to terminate conscious existence in the way ordinary production physicalism predicts.

Life after death therefore exists for a person if, after irreversible biological death, there remains or later arises conscious experience bearing a nontrivial continuity relation to the pre-death life. That relation may consist in memory, perspective, personality, dispositions, accumulated learning, an enduring subject of experience, or meaningful preservation within a larger consciousness. Strict numerical identity is not required by this operational criterion; whether it is required for personal survival remains a substantive philosophical dispute.

The definition includes:

- a surviving separable subject or soul;
- reincarnation with explicit autobiographical memory;
- reincarnation without explicit memory when genuine subject-level or psychologically meaningful continuity remains;

- participatory simulation in which an external subject continues after the death of the avatar it inhabited;
- reintegration into a larger consciousness if the prior life remains present, accessible, remembered, or constitutive of that larger consciousness.

It excludes:

- mere persistence of matter or genes;
- cultural influence or remembered reputation;
- a later reconstruction from public records that merely resembles the person, without a demonstrated continuation relation;
- the bare persistence of a cosmic mind if nothing of the prior experiential stream remains in any meaningful sense.

The definition is intentionally broader than “waking up as exactly the same autobiographical self.” If a finite conscious perspective reintegrates into a greater mind and its experiences remain genuinely present there, total conscious termination has not occurred even if the postmortem mode of identity differs from ordinary waking identity.

Memory content merely stored outside any experiencing subject does not qualify. Nor does a qualitatively identical stranger automatically qualify. The borderline is a causally connected successor: a preserved and resumed process, or an experiencing larger subject that genuinely retains the finite life. Psychological-continuity accounts may accept such successors, while stricter identity accounts may reject them. The headline uses the broader criterion; it is not a separate estimate of unchanged autobiographical identity, continuous awareness, or eternal immortality (Parfit 1984).

2. Method: Bayesian Architecture Without False Precision

The method is Bayesian in structure rather than in the sense of pretending that the inputs are measured frequencies. Every important number is a transparent judgment. The point of assigning numbers is to force hidden assumptions into the open, reveal dependence, and show which disagreements actually matter. $\text{Posterior odds} = \text{Prior odds} \times \text{Likelihood ratio}$ The model keeps route probabilities separate from evidence likelihood ratios and does not treat the mere existence of a philosophical argument as an empirical observation that can be multiplied into the odds. This staged architecture is essential to avoiding double counting and false precision.

Stage	Operation	Dependence rule
1	Allocate local architectures	State the partition and contested starting weights.
2	Update on brain dependence	Count lesions, anesthesia, and ordinary causal structure in one cluster.
3	Assign continuation conditionals	State continuity and assumed preservation mechanisms.
4	Apply empirical judgments	Use branch predictions and one log-evidence discount.

Stage	Operation	Dependence rule
5	Test sensitivity	Recalculate inputs; endpoints are not confidence limits.

The tree is a partition by the nearest local causal architecture relevant to biological death, not by ultimate metaphysics. This matters because an S1 world could have a physicalist, idealist, or dualist base reality. For survival bookkeeping, such a case is classified as S1 because the local avatar's biological death is not the death of the external participant. Only nonsimulated local worlds enter the idealist and separable/transmission branches. A simulated idealist world is assigned first to its simulation architecture, not counted again as nonsimulation idealism. This convention prevents double counting.

3. Dependence Controls and Model Discipline

Several safeguards govern the calculation. Philosophical possibilities are first represented as mutually exclusive local architectures; detailed neuroscience is then applied once; direct survival observations are assigned only to mechanisms that predict them; correlated observations receive dependence haircuts; and downstream quantities such as survival are conditioned on architecture rather than allowed to feed backward as if they were independent evidence for the ontology that generated them.

4. The Direct Empirical Survival Literature

4.1 Near-death experiences

Near-death experiences are psychologically real and often highly structured. Their phenomenology alone has limited ability to discriminate survival from an unusual brain-mediated experience. Neurophysiological and drug analogues show that several familiar features can arise without established irreversible death, but do not demonstrate an ordinary explanation for every reported observation (Timmermann et al. 2018; Martial et al. 2019, 2025).

Long’s evidence program deserves assessment beyond general phenomenology. Evidence of the Afterlife organizes nine evidence lines; his 2021 synthesis expands the argument to twelve. These include lucidity, apparent extracorporeal observations, reports from blind experiencers, experiences associated with anesthesia, life reviews, encounters with deceased persons, childhood experiences, cross-cultural similarities, and aftereffects (Long and Perry 2010; Long 2014, 2021). The lines overlap and must not be multiplied as independent evidence. NDERF’s research index includes work by other investigators as well as archive-based research.

The most discriminating category is specific information apparently acquired during an experience and subsequently corroborated. Hospital investigations cited by Long, including Sartori, Badham, and Fenwick’s patient 10 and the verified awareness account in AWARE, merit closer attention than a generic retrospective narrative. Ring and Cooper’s blind-experiencer research adds a relevant comparison: Nancy’s account had independent witness corroboration, whereas Vicki’s congenital-blindness account lacked locatable witnesses for its central observations (Ring and Cooper 1997; Sartori et al. 2006; Parnia et al. 2014). Corroboration quality varies, and accurate information does not by itself establish survival.

The published objections and responses were considered together. Cousins (2012) questions sampling and methodological transparency in Long's bestselling book; individual allegations require checking against editions and survey waves. Rush (2013) raises sensory access, chronology, memory, and interview effects in Sartori's case; Sartori (2013) disputes several proposed explanations while clarifying reporting limitations. Augustine (2022) challenges selective positive evidence in the Bigelow competition. These concerns limit confidence without disposing of each case. A merely possible ordinary explanation differs from one supported by a case's specific evidence; failure to establish an ordinary explanation is likewise not sufficient to establish a paranormal one.

AWARE II reported 567 arrests, 53 survivors, and 28 interviewed survivors; 11 interviewees described memories or perceptions suggestive of consciousness. No participant identified the concealed visual image, and one identified an auditory stimulus. These are not 567 effective visual-target opportunities: resuscitation, survival, target exposure, recall, and interviewing sharply reduce the test sample (Parnia et al. 2023). Organized EEG activity appeared during some monitored resuscitations. The recordings do not establish the exact timing or neural basis of each later recollection, and CPR can restore some cerebral blood flow. Arrest or unresponsiveness does not establish complete relevant brain inactivity or irreversible death.

Long and Woollacott (2024) document reported long-term transformations, including changes in values and fear of death. Such aftereffects strengthen the psychological significance of NDEs, but changes in belief are not independent verification of postmortem existence. Retrospective archive accounts also require checks on contemporaneous notes, independent witnesses, medical records, prior knowledge, ordinary sensory access, reporting delay, and selection.

We therefore separate corroborated perceptions from phenomenology and prospective target failures. The central model assigns the perception cluster a raw comparative factor of 1.20 to each alternative branch relative to production. This equal allocation is a simplifying, conservative judgment, not a measured likelihood ratio or a claim that all three architectures predict the observations equally well. Detailed branch-specific calibration is unavailable. The accounts are treated as one correlated cluster; the existing log-evidence discount remains. Appendix C documents the review scope and case selection, while Appendix A supplies the calculation and sensitivity to this judgment.

4.2 Childhood cases with claimed past-life memories

The strongest positive empirical material remains the small subset of childhood cases in which detailed statements were documented before the putative previous personality was identified. Retrospective cases are much more vulnerable to leakage, reconstruction, selective reporting, and search-space inflation.

The paper considers published accounts of several often-cited prerecord cases - including Kemal Atasoy, Kumkum Verma, Thusitha Silva, Iranga Jayakody, Jagdish Chandra, and Sujith Jayaratne - and distinguished two questions that are too often collapsed: whether the cases contain anomalous information, and whether anomalous information specifically implies reincarnation. Even a genuine information anomaly can be modeled by living-person telepathy, clairvoyance, deceased communication, idealist nonlocal information, or error in documentation. The cumulative result was not dismissal. The prerecord subset deserves positive weight. But the appropriate evidential target is branch-specific: it favors ontologies on

which subject-level or nonlocal continuity is natural more than it favors participatory simulation, and it does not by itself establish a reincarnation mechanism.

4.3 Mediumship and laboratory psi

Sarraf, Woodley of Menie, and Tressoldi (2021) report a positive pooled mediumship result. Blinding is useful, but replication, scoring flexibility, selection, and ordinary information pathways still matter. A pooled anomaly is not equivalent to a replicated communication channel from deceased persons.

The Transparent Psi Project failed to reproduce the particular precognition effect it tested (Kekecs et al. 2023). A subsequent three-study project with 26,483 participants also failed to reproduce the original above-chance effect. It found an opposite-direction signal that was confirmed once and then failed in its next replication (Walleczek et al. 2025). This is evidence against that specific effect and against easy aggregation of weak, selected psi findings; it is not a test of every possible anomalous mechanism. Sudduth's alternative-information objection remains: even established anomalous information would not uniquely identify survival.

This creates a coherence constraint. A survival theory cannot take mediumistic hits as evidence of robust postmortem communication and then excuse every failed cipher or lock test by saying the dead may be unmotivated, amnesic, or channel-limited. The same auxiliary assumptions must govern both successes and failures.

4.4 Purpose-built survival tests

Thouless's ciphers and Stevenson's combination-lock protocol are especially valuable because success criteria were fixed in advance. No medium supplied the keys. Two Thouless ciphers were eventually solved computationally, and Stevenson's own lock remains unopened. These are real negatives for theories that predict precise, motivated, postmortem communication.

They are not powerful negatives against broad survival as such. Survival plus preserved memory plus motivation plus a communication channel plus recipient access plus precision is a conjunction. The failure of that conjunction bears strongly on communication-capable dualism, less on idealism, and hardly at all on S1, which need not predict any in-world return channel.

4.5 Terminal and paradoxical lucidity

Lucidity in severe dementia is important for neuroscience, but current studies do not establish an extracerebral memory store. Gilmore-Bykovskyi et al. (2025) validated nine lucid episodes in three participants across 539 observations. A separate conference abstract reported caregiver-perceived lucidity in 93 of 151 enrolled patients (Tollock et al. 2025). The difference in ascertainment and validation prevents treating these rates as interchangeable.

These observations show that capacities can sometimes become accessible after severe impairment. They do not yet establish that ordinary network fluctuations explain every striking episode, or that a documented episode exceeds all plausible surviving neural capacity. The neural mechanism remains an open research question.

The broader lesson remains modestly important: severe behavioral impairment does not prove that all relevant information has been destroyed. It may reflect impaired access. But this point is compatible with sophisticated physicalism and should not be multiplied as a separate survival factor.

4.6 Deathbed and shared-death experiences

Deathbed visions and shared-death reports are recurrent across hospice literature, but the data are dominated by retrospective reports and are highly vulnerable to delirium, medication, expectation, memory reconstruction, and social contagion. They provide weak positive support at most. The model does not allow emotionally vivid cases to outweigh the methodological quality of the evidence.

4.7 Empirical conclusion

Once positive and negative observations are assigned to relevant branches and dependence is discounted, the direct survival record provides a modest rather than decisive positive update. The corroborated-perception allowance addresses information claims separately from phenomenology. Alternative information mechanisms, including those not requiring survival, remain live competitors. The net empirical change is small relative to the effects of architecture allocation and participatory frequency.

Empirical cluster	Participatory	Idealist interface	Separable / transmission
Prerecorded past-life subset	1.10	1.30	1.50
Mediumship / psi after replication discount	1.02	1.10	1.15
NDE phenomenology after neural analogues	0.97	0.92	0.85
Corroborated NDE perceptions	1.20	1.20	1.20
Prospective visual-target nulls	0.97	0.92	0.85
Cipher / lock nulls	0.95	0.90	0.70
Lucid episodes	1.01	1.02	1.03

Production-family factors are 1.00 by reference convention. These are subjective comparative judgments, not empirical effect estimates. The 0.75 exponent retains 75% of log evidence as a heuristic dependence discount. Product values and precise posterior arithmetic are retained in Appendix A. Keeping phenomenology and target-test factors fixed makes the added perception allowance auditable; it does not establish that each judgment is uniquely justified.

5. Structural Evidence: The Brain, Anesthesia, and Causal Closure

5.1 Detailed brain dependence

The strongest general evidence for production physicalism is not any one lesion or drug but the extraordinary granularity with which mental function tracks neural organization. Focal lesions alter language, recognition, personality, memory, inhibition, and value. Electrical and pharmacological interventions alter experience. Dementia progressively changes autobiographical and dispositional structure. Split-brain surgery alters information access. Developmental assembly of the brain tracks the emergence of cognitive capacities.

Production physicalism predicts this directly. Naive separable dualism has the greatest difficulty because almost every psychologically identifying capacity appears manipulable

through the brain. Yet the evidence does not discriminate equally against all non-production models. In a participatory simulation, the simulated brain can be the mandatory interface through which the external subject manifests. In interface idealism, brain events can be the public representation of deeper mental dynamics. The price those models pay is not impossibility but auxiliary structure: the tight interface mapping must be part of the theory.

The central model therefore gives brain dependence a substantial penalty against separable/transmission dualism while treating S1 and aggregate idealism more gently because those architectures can encode a mandatory interface mapping. Central structural likelihoods relative to the local-production reference are approximately $S1 = 0.80$, aggregate idealism = 0.68, and separable/transmission dualism = 0.18.

5.2 General anesthesia

General anesthesia initially looked like an independent objection to S1: if the external participant remains conscious, why does anesthesia produce an apparent blank? That framing confuses persistence of the subject with continuous reportable experience through the avatar. A full-immersion interface can simply suspend the in-world stream when the avatar brain enters the anesthetized state.

The radio analogy captures the logical point: turning off the receiver interrupts the received program without proving that the broadcast source ceased to exist. The analogy is not evidence that a broadcast source exists; it shows only why anesthesia is not an independent discriminator once brain dependence has already been counted.

A weaker defense was rejected. It is not plausible to say that patients routinely remain fully conscious through surgery, experience severe pain, and merely fail to remember it. If prolonged hidden suffering were routine, one would expect a much larger postoperative signature in trauma, autonomic disturbance, or implicit distress. Isolated-forearm work does show brief connected consciousness without later recall in a minority of patients, sometimes with signaled pain, while remembered intraoperative awareness carries a high PTSD burden. The data therefore support two conclusions simultaneously: absence of later report does not prove absence of all experience, but routine prolonged amnestic surgical suffering is not supported as a general account. This is a defeasible inference, not a proof: anesthetic agents can suppress memory, stress responses, and later trauma expression. Lack of postoperative trauma cannot by itself establish absence of all experience (Linassi et al. 2018; Sanders et al. 2017; Leslie et al. 2010).

5.3 Carroll's Core Theory and the carrier problem

Sean Carroll's low-energy effective-field-theory argument is powerful against a specific family of soul mechanisms. If person-defining information is wholly encoded in the physical brain and some unknown field or particle must read that information and carry it away at death, a new interaction of sufficient strength should leave traces in well-tested low-energy physics.

The argument does not automatically extend to every survival model. The extraction framing already assumes that the person is wholly produced and stored in the physical brain before something transports it elsewhere. On a transmission model the persisting subject need not be exported from the receiver. On idealism the physical brain is not fundamental. On S1 the ultimate subject was never wholly inside the simulated brain, and nothing resembling a synaptic data dump has to cross the simulated world's causal boundary at death.

Carroll has a residual interaction point that survives this critique: if a separable soul causally interacts with the brain during life within our low-energy physics, that interaction itself faces constraints. The model therefore folds the residual cost into the dualism brain-dependence/causal-structure likelihood rather than applying a second global penalty.

5.4 Epiphenomenalism and eliminative materialism as classification tests

Eliminative materialism and illusionism strengthen the local-production family rather than creating a new survival route. Classical eliminativism holds that some ordinary psychological categories may disappear from mature neuroscience; phenomenal illusionism goes further by denying that irreducible phenomenal properties exist as introspection represents them. If a person is wholly constituted by the physical or functional process, biological destruction terminates that process unless some independent copying, preservation, or simulation mechanism intervenes.

Epiphenomenal property dualism makes the complementary point that nonphysical consciousness does not by itself imply survival. A theory can hold that phenomenal properties are irreducible yet caused moment by moment by brain states and causally inert with respect to the physical world. On that architecture, irreversible brain failure can end the phenomenal stream even though physicalism is false. The model therefore reserves the high survival conditional for separable or transmission dualism, not for dualism indiscriminately.

For bookkeeping, these brain-dependent termination variants are included in the residual local-production family. This improves the taxonomy without changing the central 47% result: the 0.350 starting mass is a local-production category with no specified ordinary continuation mechanism, not a claim that every member is metaphysically physicalist.

6. Participatory Simulation: The Central Structural Alternative

6.1 S1 and S2

Two simulation architectures matter. In generated simulation, the subject is produced within the implementation. Its continuation depends on whether the conscious process is preserved, resumed, or meaningfully reinstated. Death of a simulated biological body need not mean deletion of the implementation; an operator might retain its state. Generated simulation therefore does not imply either survival or termination by itself.

In participatory simulation, a pre-existing subject experiences an avatar through an interface. Avatar death need not kill that subject, but the interface could alter memory, the participant could also die, or the avatar's perspective could fail to remain meaningfully available. The high continuation conditional used below is an explicit assumption about those possibilities, not part of the definition of simulation.

For compatibility with the historical notation, S1 means participatory simulation and S2 means generated simulation. The prose uses the descriptive names where possible.

6.2 The simulation/Descartes interaction architecture

A recurring objection to nonphysical mind is Princess Elisabeth's question to Descartes: how can an unextended mind cause motion in extended matter? Participatory simulation supplies a

concrete local architecture that dissolves the problem at the simulated level. external intention → simulator/interface → ordinary in-world neural precursor → in-world action An in-world neuroscientist can observe an apparently complete sequence of neural causes and still be wrong to infer that no external subject exists. The simulator maps the participant's intention into an ordinary simulated neural state, after which all downstream events obey local law. Apparent local causal completeness can therefore coexist with a deeper generating system. This needs a qualification. If a complete local state plus fixed local laws uniquely determines the next state, an independently variable external input cannot change that next state while leaving those conditions fixed. The interface must either be represented within the local transition rule, generate a whole internally consistent history, or produce interventions that an incomplete local description fails to detect. A chain of ordinary-looking neural precursors alone does not prove strict local causal closure. The counterexample challenges an inference from observed completeness to ultimate completeness; it does not show that arbitrary external control is compatible with an autonomous deterministic local dynamics.

If the base-level participant is itself physically instantiated, there is no ultimate Cartesian interaction problem: a physical participant interfaces with a physical simulator, while relative to the simulated world the mind is outside local physics. If the base-level participant is genuinely nonphysical, Elisabeth's objection is relocated to base reality rather than eliminated. This distinction shows why local dualism does not entail ultimate dualism.

Participatory simulation also clarifies two standard objections to epiphenomenalism. If consciousness is the way an external participant experiences the simulation, the subject can have genuine intentions and phenomenal judgments while the simulator translates those states into ordinary in-world neural precursors. The in-world causal chain can remain completely lawful, yet reports about pain or color can track the subject's experience through the deeper interface architecture. This is a strong reply to the local mental-causation and phenomenal-judgment problems.

The architecture is not strict global epiphenomenalism, however. If the external subject is physically instantiated, the base-level interaction can be entirely physical; if a genuinely nonphysical intention changes the interface state, consciousness is causally efficacious at the deeper level. The point is therefore not that strict epiphenomenalism predicts survival, but that local physical closure can coexist with genuine subject-level causation and post-avatar persistence under S1.

6.3 Strong computational functionalism

Bostrom's original observer-counting argument assumes substrate independence: a computer running the right program must be capable of hosting conscious minds. The project therefore had to assess strong computational functionalism rather than silently assume it.

The working credence in strong computational functionalism is roughly 25 percent. This is an elicited philosophical judgment, not a probability estimated from current AI experiments. This is not because current AI has become intelligent without obvious consciousness; intelligence and complexity alone are not the functionalist thesis. The downward pressure comes from the metaphysical consequences of organizational invariance and from the inadequacy of several common arguments that purport to force substrate independence.

A new 2026 result deserves an explicit update to that discussion. Anthropic's interpretability work identifies a small, privileged set of internal representations in Claude - J-space - whose

contents can be reported, deliberately summoned and held, used in reasoning, capacity-limited, and broadly broadcast to downstream computation. Gurnee and colleagues argue that this has the functional hallmarks of a global workspace. Bostrom and Shulman separately treat substrate independence and global-workspace-like theories as serious candidates for digital mentality.

The result is evidence of architectural convergence, not a consciousness meter. It supports the proposition that access-like organization can emerge in nonbiological computation, but it does not establish that there is something it is like to instantiate that organization. Anthropic makes the same access/phenomenality distinction, and McClelland's 2026 argument for agnosticism about artificial consciousness explains why cross-substrate extrapolation remains contested. The paper therefore treats J-space as modest upward pressure on strong computational functionalism, but not as a quantified likelihood ratio large enough to re-round the 25 percent central working weight on current evidence.

That restraint matters for the survival calculation. If future evidence moves the functionalism weight materially upward, probability shifts from S1 toward S2 within simulated worlds. That can actually lower postmortem survival on this model, because an S2 subject is generated by the in-world computation and need not persist when that process ends, whereas an S1 subject exists outside the avatar. The direction of the update is therefore not 'more AI consciousness implies more survival'; it depends on which simulation architecture the consciousness evidence supports.

6.4 The hydraulic computer

A sufficiently large system of pipes, valves, and flowing water could in principle instantiate the same abstract computation as an electronic system. Strong organizational invariance is committed to saying that if the causal topology were exact, the hydraulic system would instantiate the same pain, color experience, nostalgia, and unified perspective as the electronic implementation.

The lineage runs from Leibniz's mill through Block's population-of-China objection and Searle's water-pipe brain simulator. The formulation developed in this project focuses specifically on phenomenal properties and on the illicit intuitive work performed by electricity and silicon: once the computation is translated into plumbing, the claim that causal organization alone generates consciousness becomes easier to see as a substantive metaphysical commitment rather than an obvious consequence of computation.

The functionalist reply deserves equal clarity: changing silicon to plumbing does not alter the stipulated causal organization, so an intuition about unfamiliar materials is not independent evidence against organization-based consciousness. The thought experiment exposes the commitment. It does not refute it.

6.5 Gradual replacement and threshold death

The gradual-neuron-replacement argument asks where consciousness could disappear if biological neurons are replaced one at a time by functionally identical artificial units. The standard functionalist intuition is that no individual replacement seems capable of extinguishing the whole field of experience, so the fully replaced system must remain conscious.

The reply developed here is threshold death. Functional continuity can outlast the biological conscious subject. At some substrate-dependent threshold the subject may cease even though the artificial system continues the person's functional profile. Difficulty locating the threshold does not show there is none. Thresholds and sharp transitions occur in physical systems, although an exact discontinuity in a finite brain would need its own mechanism. This does not prove consciousness has a critical substrate parameter; it shows that 'one neuron cannot make the difference' is not a logical refutation.

6.6 Dancing qualia and the anesthesia toggle

Chalmers's dancing-qualia argument says that functionally identical biological and silicon realizations with radically different qualia would permit enormous phenomenal changes that the system could not notice or report. That is supposed to be absurd. The strongest project reply targets absent qualia. Imagine toggling control from the conscious biological system to a functionally equivalent nonconscious surrogate and later back. During the surrogate interval the system continues speaking, discriminating, and maintaining functional memory, but there may be no phenomenal subject. When the biological system returns, the subject resumes with the functional memory state supplied by the system. The report that 'nothing changed' is generated by function, not by a subject who phenomenally witnessed the interval.

The objection then shifts from contradiction to the reliability of functional self-ascription. If introspective reports are functionally generated while phenomenality depends on additional substrate properties, reports of continuity need not guarantee phenomenal continuity. The argument therefore remains powerful as an intuition pump but not as a deductive proof of substrate independence.

This reply has a cost. If phenomenal changes can occur without any functional ability to notice them, the theory owes an account of why ordinary introspective reports reliably track experience at all. Actual anesthesia does not settle the thought experiment, because it changes brain function; the stipulated surrogate holds the relevant organization fixed. Neither side can treat that empirical analogy as a direct test of the stipulated case.

7. Bostrom Reframed: From Generated Minds to Participatory Observer Lives

7.1 Functionalism and the generated-observer argument

Bostrom's original trilemma relies on computed minds being conscious. If computational organization is insufficient for consciousness in the relevant implementations, counting software runs does not count conscious observers. Denial of strong functionalism alone is weaker: particular nonbiological, quantum, embodied, or substrate-dependent implementations might still be conscious. The ordinary counting argument therefore requires a justified consciousness premise for the implementations being counted. J-space makes the computational route somewhat more credible at the level of access architecture, but does not by itself establish the phenomenal premise that observer counting requires.

7.2 What recoupling does not show

It does not follow that simulation itself becomes unlikely. The simulation hypothesis is broader than Bostrom's pure generated-mind implementation. If a technologically advanced system can

construct convincing environments but cannot manufacture phenomenality computationally, it can still populate those environments with pre-existing conscious subjects.

The relevant anthropic unit then changes. Instead of asking how many simulated minds are generated, one asks how many immersive conscious lives are experienced by each genuine subject. For one comparable base-life equivalent per subject and k comparable simulated-life equivalents, the simulated fraction is $k / (1 + k)$. Thus $k = 1$ already yields one simulated life for every base life; $k = 4$ yields 80 percent of life-histories simulated; $k = 9$ yields 90 percent. No civilization needs to manufacture trillions of new conscious programs. Repeated lives per persistent subject can recreate selection pressure under a life-equivalent measure. Unequal duration, unequal relevance to our evidence, many base-life equivalents, and memories of earlier immersions all change the denominator. Counting persons gives a different result from counting experienced lives or moments.

7.3 A modified participatory trilemma

The participatory version distinguishes three alternatives:

- Advanced civilizations almost never develop full-immersion interfaces connecting genuine conscious subjects to simulated worlds.
- Civilizations with such interfaces almost never place subjects into enough comparable immersive lives to dominate the relevant observer measure.
- Otherwise, a substantial, potentially dominant fraction of relevant conscious lives are lived inside simulations.

This is not Bostrom's original theorem. The third conclusion requires enough relevant lives per base-life equivalent, plus a justified sampling measure. Participation makes such an argument available under some substrate-dependent theories; it does not establish its empirical premises.

7.4 Participatory frequency analysis

The model avoids the crude question 'would posthumans enjoy playing Earth?' and decomposes participatory frequency into technological, behavioral, duration, and reference-class factors. Let R be the ratio of relevant S1 observer-life equivalents to relevant base-life equivalents. A useful schematic is: $R \approx (e_H + e_A) \times d \times m$ where e_H is the expected number of voluntarily chosen immersive lives per external subject across the population, e_A is the expected number of AI- or institution-driven immersive lives per subject, d is relative experiential duration, and m is the fraction of immersive lives that belong to the reference class relevant to our own apparently ordinary human existence.

The analysis yields a wide range. Low uptake, short immersions, or a narrow reference class can drive the implied S1 frequency into the low teens; widespread repeated immersion can push it above one half. The central input is kept near 0.30 for $P(S1 \mid \text{not-F})$, with a serious range roughly 0.15 to 0.50 rather than treating any single value as a measured derivation.

Here not-F is a shorthand complement of the strong functionalist thesis, not a homogeneous theory that makes all generated minds unconscious. The frequency input already averages over further theories and technological possibilities. A more complete analysis would divide that complement into biological, quantum, other substrate-dependent, and consciousness-first

views before estimating the simulation shares.

7.5 AI as simulator and experiential data as motive

A further project development reduces the anthropomorphic character of the frequency objection. The simulator need not be a posthuman individual who personally wants to 'play' a painful eighty-year life. It may be an advanced AI or AI-governed system. The conscious participant and the simulator/operator are distinct roles. An operator could generate an environment, connect a genuine participant, and collect information from the resulting experiential trajectory. Modern machine learning makes the instrumental value of data obvious. An advanced intelligence could have reasons to generate enormous numbers of richly structured trajectories for prediction, planning, psychology, medicine, social modeling, alignment, strategy, or the study of conscious agents themselves. In present terminology this resembles synthetic-data generation, except that the trajectories can be generated by genuine conscious participants rather than by nonplayer characters implemented computationally.

The strongest objection is equally clear. If nonconscious computational agents produce all the useful behavioral data, genuine participants may add little. The AI-data motive therefore becomes strongest if phenomenal consciousness makes a causal difference to trajectories, if some information cannot be duplicated by zombie computation, or if consciousness itself is the object of study. The model treats the AI/data argument as a genuine positive driver, not a decisive one.

8. Idealism Split into Survival-Relevant Families

A single 'idealism = 35% survival' conditional concealed architectures with very different implications. The model separates three families.

8.1 Dissociative analytic idealism

On Kastrup-style analytic idealism, individual minds are dissociated alters of a universal field of consciousness. Biological death plausibly corresponds to the dissolution of the dissociative boundary. Under a strict autobiographical-identity standard this is ambiguous. Under the definition used here, however, reintegration counts as survival if the finite life's experiences remain genuinely present or accessible within mind-at-large.

The survival conditional is therefore high but not near certainty: the universal consciousness persists by hypothesis, yet the degree to which the former finite perspective remains structured and accessible is not fully specified.

8.2 Conscious-agent/interface idealism

Hoffman-style conscious-agent theories are more directly survival-friendly. Spacetime and brains are interface representations rather than ultimate ontology, while conscious agents are fundamental. The disappearance of the body icon therefore does not straightforwardly imply disappearance of the agent. The uncertainty concerns the mapping between the ordinary human self and the deeper agent network, not whether consciousness as such depends on the biological icon.

8.3 Cosmopsychism and impersonal objective idealism

Cosmopsychism makes a cosmic or universal conscious reality fundamental while treating local human subjects as derivative aspects, localizations, or organizations. This does not automatically imply personal immortality. A local center might dissipate like a wave in an ocean. Yet under the broad survival definition, reintegration counts if the prior experiential content remains in the encompassing consciousness. Cosmopsychism is therefore less survival-friendly than interface idealism but more favorable than a simple undifferentiated idealist treatment would imply.

Idealist family	Internal weight	Illustrative continuation
Dissociative analytic idealism	35%	0.70
Conscious-agent / interface idealism	30%	0.80
Cosmopsychism / objective idealism	35%	0.55
Weighted average	100%	0.6775, represented as 0.68

These are consequence assumptions for nonsimulation idealist branches, not measured survival rates. Cosmopsychism need not be idealist in every taxonomy; it belongs here only insofar as it gives the relevant consciousness-first local architecture.

The aggregate is rounded to 0.68 in the worked model. Fundamental consciousness alone does not guarantee continuation of a finite perspective. A stricter criterion for personal identity would lower these conditionals, and possibly change the headline substantially.

9. Solipsism as a Subject-Multiplicity Parameter

Strict one-subject solipsism should not be added as a fifth top-level ontology. Once made sophisticated enough to explain the stable world, it becomes either a one-participant version of S1 or a one-local-subject version of consciousness-primary reality. The cleaner parameter is N, the number of genuine finite phenomenal subjects associated with this world. $N = 1 \rightarrow$ single-subject/solipsistic branch; $N > 1 \rightarrow$ multiplayer S1 or plural idealism

9.1 Evidence against $N = 1$

The strongest discriminator is the apparent autonomy of other minds. Other people know facts we do not know, make discoveries we did not anticipate, conceal information, defeat us intellectually, maintain long unobserved histories, and later produce independently verifiable records. Multiplayer S1 and plural idealism predict this naturally. A one-player simulator can generate the same behavior, but it must supply the hidden cognitive machinery.

Modern AI makes the inference less automatic only conditional on the relevant AI systems lacking phenomenal consciousness. That premise remains unsettled. If highly autonomous behavior can be produced without phenomenal consciousness, then surprise, intelligence, strategic concealment, and even intellectual superiority are less decisive evidence of another phenomenal subject than they once appeared. For other humans, shared biology, development, vulnerability, and neural organization remain powerful evidence alongside behavior. Uncertainty about AI consciousness does not erase that broader case.

Persistent hidden information creates a different pressure. A sealed random number can remain stable for decades without anyone observing it. A solipsistic generator can preserve the information trivially, but once it does so the theory has admitted a vast persistent information-bearing structure independent of the sole subject's current experience. It has economized on other consciousnesses, not necessarily on external ontology.

9.2 The problem of evil reverses direction under benevolent single-subject generation

A single-subject world radically alters the amount of actual suffering. Under ordinary plural realism, reports of war, animal predation, disease, torture, and prehistoric suffering correspond to enormous numbers of phenomenal victims. Under $N = 1$, only the sole subject's first-person suffering is phenomenally instantiated; nearly all third-person suffering is representational.

apparent suffering of others $\neq$ actual phenomenal suffering, if $N = 1$

This does not make solipsism true. It changes the likelihood under a benevolent-generator hypothesis. A world containing immense apparent suffering but comparatively tolerable first-person suffering is more compatible with benevolent single-subject generation than with a benevolent world containing all the apparent sufferers as genuine conscious victims. Conversely, severe gratuitous first-person suffering is a direct negative for benevolent solipsism.

The effect must be counted only once. It belongs inside $P(N = 1 \mid \text{benevolent generator, observed suffering})$. It should not separately raise creator probability, solipsism probability, and survival probability.

9.3 Subject-multiplicity estimates

The central values are roughly $P(N = 1 \mid S1) = 0.20$ and $P(N = 1 \mid \text{idealism}) = 0.09$, with wide uncertainty. In the normalized model this places single-subject reality in roughly the mid-single-digit to high-single-digit percentage range, most of it inside one-player S1 rather than objective idealism.

Solipsism does not add new survival probability mass. A specified theory might assign different continuation conditionals to one-subject and many-subject versions. The present calculation uses an aggregate within each architecture; it does not add a separate numerical increase for singleness. Subject count also changes the problem-of-evil analysis. Treating it as a separate 5-10 percent top-level route would be double counting.

10. The Bayesian Model

The four branches are bookkeeping classes ordered by the nearest architecture relevant to bodily death. They are not four mutually exclusive ultimate ontologies. The local-production residual contains generated simulation and nonsimulation production theories, including some brain-dependent nonphysical views. A small 0.02 continuation conditional represents unspecified rescue or preservation mechanisms; it is not an empirical frequency.

The 0.18 idealist and 0.22 separable/transmission starting weights are unconditional masses restricted to nonsimulation branches. They are not inferred from the companion ontology paper. Together with the simulation inputs, they define a representative prior that sums to one. They do not reconstruct a full joint distribution over consciousness theory and ultimate

ontology.

Input	Value	Meaning
Strong computational functionalism	0.25	Working philosophical weight
Simulation given strong functionalism	0.40	Technology, motives, and observer selection
Participatory share of functionalist simulations	0.25	Remaining simulated mass is generated subjects
Participatory simulation outside strong functionalism	0.30	Aggregate nonfunctionalist judgment
Nonsimulation idealist interface	0.18	Unconditional starting mass
Nonsimulation separable / transmission	0.22	Unconditional starting mass

The simulation arithmetic is $0.25 \times 0.40 \times 0.25 = 0.025$ for participatory simulation with functionalism, and $0.75 \times 0.30 = 0.225$ outside it. Total participatory mass is therefore 0.250. Generated mass is $0.25 \times 0.40 \times 0.75 = 0.075$. The nonsimulation production residual is 0.275; combining it with generated simulation gives the 0.350 local-production family.

Architecture	Start	Brain factor	Structural weight
Local production	0.350	1.00	49%
Participatory simulation	0.250	0.80	28%
Idealist interface	0.180	0.68	17%
Separable / transmission	0.220	0.18	6%

Architecture	Continuation conditional	Empirically updated weight
Local production	0.02	45%
Participatory simulation	0.92	30%
Idealist interface	0.68	20%
Separable / transmission	0.90	5%

Output percentages are rounded for circulation. Exact normalized weights and unnormalized masses appear in Appendix A. The displayed whole percentages are not claims of precision of belief.

Using the unrounded structural weights gives about 44% continuation. Applying the discounted empirical factors from Section 4.7 and renormalizing gives about 47%. Appendix A retains the exact arithmetic.

The computation uses the law of total probability: sum each updated architecture probability multiplied by its continuation conditional. It holds the continuation conditional fixed within each architecture when applying empirical evidence. That is a simplifying assumption: some communication evidence would also update survival versus nonsurvival subbranches within a single architecture. A fully specified model would calculate those subbranches separately. The

modest empirical shift should therefore be read as the result of this illustrative scoring convention, not a uniquely determined verdict of the literature.

Changed-input calculations show which assumptions carry the estimate.

11. Sensitivity and Dependency Analysis

The headline estimate is subjected to deliberately skeptical sensitivity tests so that the reported probability does not outrun the argument. The principal vulnerabilities are prior architecture, participatory frequency, anthropic measure, brain-dependence penalties, and the continuity criterion under idealism.

11.1 Starting ontology allocation is load-bearing

Nearly every model of this kind is vulnerable to prior architecture. If idealism and dualism are both assigned much smaller pre-neuroscience weights, the posterior falls substantially. This is not a defect unique to survival analysis; it is the unavoidable consequence of comparing ontologies with different survival implications. The paper therefore does not call Stage 1 a uniquely neutral prior.

A skeptical allocation closer to idealism 0.10 and dualism 0.15 can push the result into the mid-30s or below depending on the S1 setting. Conversely, more generous consciousness-first starting weights easily lift the posterior above one half. The paper therefore reports ranges rather than treating 47 percent as an objective frequency.

11.2 Participatory frequency is the largest single live variable

The observer-life argument is structurally important but contains dependent auxiliaries: full-immersion technology must arise; interfacing genuine consciousness must be possible; repeated lives must be used at scale; durations must be substantial; a meaningful fraction of those lives must belong to our reference class; and the anthropic reference measure must be legitimate.

The AI/data argument strengthens the motive side but does not eliminate the zombie-data objection. The central value of 0.30 for $P(S1 \mid \text{not-F})$ is therefore a judgment after argument, not a derivation. A value near 0.15 materially lowers the headline; a value near 0.50 raises it dramatically.

11.3 Anthropics remains technically unsettled

Different ways of averaging over uncertainty about S1-capable worlds can produce different observer measures. For example, $q \cdot R / (1 + R)$ and $qR / (1 + qR)$ are not identical. For the same q between zero and one and positive R , $qR / (1 + qR)$ is at least $qR / (1 + R)$. The direction is fixed in that matched comparison. What remains contestable is which conditioning and sampling assumptions justify either expression. The model therefore treats anthropic arithmetic as a family of models rather than pretending one expression is mandatory.

11.4 Brain dependence can be scored more harshly

A skeptical model can reasonably say that S1 and idealism are being designed to fit the very brain dependence they are supposed to predict. The central model therefore lowers S1's brain LR to 0.80 and aggregate idealism to about 0.68; harsher values remain plausible and are

included in the uncertainty envelope.

11.5 Idealist continuity remains partly definitional

Once reintegration counts, idealism becomes more survival-friendly, but the mere persistence of universal consciousness does not guarantee meaningful preservation of the finite perspective. The aggregate survival conditional of roughly 0.68 is therefore deliberately below the most survival-friendly formulations.

11.6 Direct paranormal evidence has little leverage

Replacing the entire direct empirical update with mildly negative factors gives about 41%, compared with the central 47%. Setting only the positive NDE-perception row to neutral gives about 44%. The empirical literature matters, but architecture assumptions remain more influential than this modest evidential allowance.

11.7 Creator residual removed from headline

A residual creator-instituted continuation route is set to zero in the headline model because creator hypotheses overlap heavily with idealism, S1, dualism, and the problem of evil. Creator-instituted survival can be shown in sensitivity analysis rather than silently added.

11.8 Solipsism and the problem of evil cannot be double counted

The problem-of-evil advantage of single-subject reality is a conditional interaction, not a new route. It raises $P(N=1)$ inside benevolent-generator branches if the observed distribution of apparent suffering fits that architecture. It must not be multiplied again as independent evidence for survival.

11.9 Reproducible sensitivity calculations

Each scenario below uses the same formula and empirical discount as the central case. Unless specified, every other input stays at its central value. These are alternative assumptions, not probabilities assigned to the scenarios themselves.

Scenario	Changed input	Continuation
Central	All declared inputs	47%
Lower participatory frequency	$P(\text{participatory} \mid \text{not-F}) = 0.15$	34%
Higher participatory frequency	$P(\text{participatory} \mid \text{not-F}) = 0.50$	64%
Lower nonsimulation alternatives	Idealist start 0.10; separable start 0.15	36%
Combined skeptical inputs	Frequency 0.15; idealist 0.10; separable 0.15	25%
Stricter idealist continuity	Idealist continuation conditional 0.35	40%
Mildly negative empirical factors	Replace final empirical factors of all alternative branches with 0.90	41%
Neutral corroborated-perception row	Set only the positive NDE perception factors to 1.00	44%

Scenario	Changed input	Continuation
Positive NDE benchmark	Set only the positive NDE perception factors to 1.10	45%
Stronger positive NDE allowance	Set only the positive NDE perception factors to 1.30	48%

The 20-70% envelope is a broad working sensitivity region, not a guaranteed range across reasonable priors. Combining participatory frequency 0.12, idealist start 0.10, separable start 0.15, brain factors 0.55 / 0.45 / 0.10, idealist continuation 0.45, and final alternative empirical factors 0.90 yields about 13%. Assigning very low weight to every speculative continuation route can lower it further.

Conversely, participatory frequency 0.50, brain factors 0.90 / 0.80 / 0.30, participatory continuation 0.95, idealist continuation 0.80, and final empirical factors 1.05 / 1.15 / 1.00 yields about 69%. These examples replace the whole empirical factor when specified, rather than adding another row. They expose the cost of stacking assumptions in either direction.

12. What the 47 Percent Means - and Does Not Mean

The central estimate is not a 47 percent measured chance that a particular person survives. It is the output of an explicit comparative model built from disputed ontological judgments. Its value lies in showing what assumptions make termination nearly certain and what assumptions make survival more likely than not.

The most striking result is negative: modern brain science does not by itself warrant the common inference from 'mental life depends tightly on the brain' to 'the brain must produce the ultimate subject and death therefore annihilates it.' That inference is natural under production physicalism, but S1 and some idealist interface architectures predict most of the same dependence.

The second striking result is that the direct paranormal literature is not carrying the case. One can be skeptical of NDE veridicality, mediumship, psi, terminal lucidity, and reincarnation claims and still assign substantial probability to postmortem continuation if one gives serious weight to participatory simulation or consciousness-first ontology.

The third result is that simulation and functionalism should not be conflated. Bostrom's generated-mind argument requires consciousness in the implementations being counted; participatory simulation instead requires genuine subjects and a workable interface. Rejecting strong functionalism does not eliminate every conscious artificial implementation, but it changes which populations may legitimately enter the count.

The low-preservation benchmark is below 50 percent; the expanded analysis in Appendix D does not establish that side of one half robustly. Generated-mind continuation above about 37.5% crosses the threshold with other inputs fixed. Neither that conditional nor a new preferred headline is established by the interview. More skeptical architecture weights can still produce much lower estimates.

13. What Would Move the Estimate Most

- Participatory frequency: better arguments about how often advanced systems would place genuine subjects into immersive lives, especially whether conscious trajectories provide

information unavailable from nonconscious computation.

- **Substrate dependence:** evidence about which biological or other physical properties phenomenality requires would change the balance between generated and participatory simulation.
- **Architecture priors:** a more principled starting allocation among local production, participatory simulation, idealist interfaces, and separable/transmission theories would address a major source of disagreement.
- **Idealist identity:** a clearer account of whether dissociation and reintegration preserve meaningful experiential structure would improve the continuation conditional.
- **Past-life documentation:** a forensic scorecard should assess leakage, chronology, incorrect statements, search-space size, and independent records case by case.
- **Communication tests:** a preregistered cryptographic test with a sufficiently large keyspace and controlled access could yield a discriminating hit or null for specified communication-capable models.
- **Simulation signatures:** reproducible evidence of code, external intervention, resource constraints, or identifiable operators could favor particular simulation architectures. It would not alone establish conscious participation or survival.
- **Shared anomalous information:** independent reports with ordinary information channels excluded could discriminate specified models. Shared reports alone would not refute a sufficiently flexible one-subject generator.

About the Author

Rodney Malpert holds a Juris Doctor and a Master of Arts from Cornell University and completed all Ph.D. requirements in Government except the dissertation. His Cornell graduate work included substantial coursework in analytic philosophy. He practiced law for many years in Silicon Valley and taught law as an adjunct professor at Southern Methodist University and Arizona State University. He is a United States Chess Federation Senior Master. His current research uses human-AI collaboration to examine philosophy of mind, probability, consciousness, identity, and possible postmortem continuation.

Website and current papers: RodMalpert.com

14. Intellectual Contributions and Attribution

This project was intentionally collaborative and adversarial. Attribution below identifies contributions made within the project. Unless explicitly stated, no claim of historical priority is made.

Rodney Malpert originated the project's broad operational definition of survival and repeatedly insisted that neither historical religious assumptions nor contemporary physicalist defaults be treated as automatically privileged. He introduced and developed participatory simulation as a survival route and pressed the distinction between local and ultimate ontology.

Malpert independently formulated the simulation/Descartes interaction architecture in which an external intention is rendered by the simulator as an ordinary in-world neural precursor, permitting deeper mental causation to coexist with apparent local causal closure. He also

developed within this project the hydraulic-computer reductio against strong computational substrate independence; the threshold-death response to gradual neuronal replacement; and the anesthesia-toggle response to Chalmers's absent, fading, and dancing qualia.

Malpert further developed the critique that Carroll's carrier formulation presupposes an extraction picture when generalized beyond its proper scope; the radio-turned-off argument that anesthesia has no independent force beyond the brain-dependence question; and the postoperative-trauma objection to the claim that patients remain routinely conscious through surgery but merely forget it.

Two further contributions are central to the simulation analysis. First, Malpert reframed the Bostrom-style selection effect for participatory simulation: if computed minds are not conscious, the relevant multiplier can be the number of immersive lives experienced by pre-existing conscious subjects rather than the number of generated minds. Second, he introduced the possibility that advanced AI rather than posthuman individuals may operate the simulations and may have an instrumental demand for large quantities of rich experiential data, providing a non-anthropomorphic motive for high participatory frequency.

Malpert also independently identified the problem-of-evil reversal under single-subject reality: if only one finite phenomenal subject exists, almost all apparent third-person suffering is representational rather than actual phenomenal suffering. This does not establish solipsism, but it changes the likelihood of observed suffering under benevolent-generator hypotheses and creates a distinctive contrast between apparent remote suffering and first-person suffering. Malpert additionally required explicit evaluation of epiphenomenal property dualism and eliminative materialism/illusionism so that neither nonphysicalism nor physicalism was treated as a single undifferentiated family. He proposed that the standard epiphenomenalist objections from mental causation and phenomenal judgment may be substantially mitigated under participatory simulation when consciousness is the external participant's mode of experience and the simulator maps the participant's intentions and judgments into lawful in-world neural precursors. The paper adopts that proposal as a coherent S1 architecture while noting that it is not strict global epiphenomenalism if consciousness is causally efficacious at the deeper level.

A further September 2026 intervention by Malpert was to connect Bostrom's computationalist/global-workspace discussion and Anthropic's J-space result directly to the S1/S2 architecture. He pressed both sides of the update: convergence toward a global-workspace-like organization is evidence that nonbiological systems can reproduce access-relevant structure, but it must not be conflated with proof of phenomenality. He also identified the counterintuitive survival implication that stronger S2 plausibility can shift probability away from the survival-favorable S1 branch. The evidential calibration and present prose were developed collaboratively.

ChatGPT developed the initial empirical taxonomy and H0/H1/H2 distinction; built the brain-dependence comparison; integrated simulation, idealism, dualist, creator, epiphenomenal, and eliminativist classifications; and translated the conceptual arguments into the conditional S1/S2, functionalism, participatory-frequency, idealist-subfamily, solipsism, and sensitivity models.

Claude supplied major adversarial structural corrections involving explicit staging, purpose-built negative survival tests, sensitivity analysis, recoupling of Bostrom's generated-mind argument to substrate independence, and branch-specific empirical likelihoods. Grok served as an additional external critic, particularly on total simulation

probability, participatory frequency, brain-dependence penalties, and dependence control.

The September 30 review by ChatGPT (GPT-6.1 Sol) recalculated the numerical model, made the empirical discount explicit, tightened identity and causal-closure claims, checked key primary sources, corrected reference metadata, and rebuilt the layout. Claude and Grok are credited for their earlier contributions; they did not independently repeat this latest audit.

Following feedback from Dr. Jeffrey Long, Malpert directed a review of Long's published evidence, NDERF accounts, hospital investigations, and substantive published criticisms and replies rather than asking Long to nominate his strongest cases. After discussion with ChatGPT, Malpert ultimately judged that the modest positive NDE allowance discussed at 1.10 should be increased to 1.20. This records the author's evidential judgment, not a measured effect. Because phenomenology, positive corroborated perceptions, and prospective failures have different evidential roles, the numerical implementation is the distinct corroborated-perception row in Section 4.7. The conversational 1.10 benchmark should not be confused with an existing uniform NDE multiplier in the model. ChatGPT organized the source review, distinguished documentation levels, drafted the synthesis, and recalculated the empirical and sensitivity tables. The final choice of evidential weight is Malpert's.

The method is adversarial rather than consensus-driven: survival-friendly inferences are stress-tested against stronger neuroscience, negative tests, dependence, and alternative mechanisms. The central 47 percent figure is therefore intentionally lower than a model that simply stacks every continuation-friendly assumption.

15. Conclusion

The strongest conventional case against life after death is formidable: personality, memory, access, and cognition depend in extraordinarily fine-grained ways on the brain. A naive separable soul theory pays a large evidential price. Purpose-built communication tests have not delivered the results survival advocates would most like to have. Near-death phenomenology has increasingly plausible neurobiological models. No serious analysis should minimize those facts.

Those facts do not logically entail annihilation under every architecture. How much probability the surviving alternatives deserve depends on their prior plausibility and explanatory costs. Participatory simulation provides a coherent architecture in which the apparent brain is an interface while local neuroscience remains completely predictive. Objective idealisms provide architectures in which brains are appearances or local organizations within a more fundamental mental reality. Neither view is established, but both survive the brain-dependence evidence much better than naive dualism.

The simulation analysis also changes when functionalism is no longer assumed. If computation alone cannot generate phenomenal consciousness, Bostrom's pure simulated-observer multiplier fails. But the simulation hypothesis does not. Conscious subjects can participate from outside the simulated world, and repeated immersive lives can restore a Bostrom-like observer-life selection effect. AI operation and experiential-data generation supply a further reason such lives might occur at scale.

Solipsism remains a minority conditional possibility rather than a preferred ontology, but it is philosophically instructive. Modern AI weakens the old inference from autonomous intelligent behavior to separate phenomenal minds, while the problem of evil changes radically if most

apparent sufferers are not phenomenally instantiated. These observations reinforce the broader methodological lesson: theories that are conceptually distinct can be empirically difficult to separate, and Bayesian comparison requires predictive specificity rather than mere compatibility.

Under the low-preservation benchmark, broad conscious continuation is approximately 47 percent. Appendix D separates generated minds from ordinary local production and exposes the preservation conditional: a 50% conditional yields 51.2% overall with other inputs fixed. The result is therefore not robustly below one half. These are model-conditioned judgments, not measured probabilities; neither the benchmark nor its extension establishes a model-independent probability floor.

16. Supplemental Analysis

The following analyses provide additional support and boundary checks but are not assigned independent multipliers when their evidential content is already represented elsewhere. They are included to make the central weights intelligible without double counting.

16.1 Classical arguments: Plato, Aristotle, Aquinas, and Descartes

The major classical arguments for the soul and afterlife are historically important but should not be converted into likelihood ratios merely because they exist. Plato's cyclical, recollection, affinity, and life-bearing arguments largely overlap modern issues such as the hard problem, unity of consciousness, and mind-body ontology. Aristotle's hylomorphism is not straightforwardly survivalist, and the arguments associated with Aquinas and Descartes gain or lose force through the same contemporary metaphysical questions already represented elsewhere in the model.

Aquinas provides a more explicit survival route: intellect is not reducible to a bodily organ, therefore the rational soul is subsistent and capable of existing apart from the body. The modern difficulty is that abstract thought, language, memory, and personality are deeply vulnerable to neural manipulation. Descartes gives the canonical real-distinction argument, but his interaction problem and the brain-dependence evidence are severe. None of these arguments was discarded; each was relocated to the ontology allocation where it belongs rather than multiplied as though it were an empirical observation.

This methodological correction was important. A philosopher's argument can alter a rational agent's credence by changing the plausibility of an ontology, but 'Plato made an argument' is not itself an event whose probability can be compared under survival and non-survival. Treating philosophical reasoning as if it were a data point would double count the same considerations that already shape the structural prior.

16.2 Detailed past-life case analysis

The childhood past-life literature was not treated as a homogeneous block. The project distinguished prerecord cases - where a child's statements were documented before a candidate previous personality was identified - from retrospective cases in which information could have been shaped by later search, family contact, memory, or selective reporting. The prerecord subset is small, but it is the subset with genuine diagnostic value.

The earlier case discussion is retained as a comparison of published reports, not as an independent forensic validation of their original records. Kemal Atasoy was reported to have made detailed statements about an Armenian Christian life, an Istanbul waterfront home, a church, wealth, and business details before identification; a substantial number were later judged confirmed, while others were partial or doubtful. Kumkum Verma's statements were written by an aunt and included locality, family, occupational, and household details later matched to a deceased woman. Thusitha Silva, Iranga Jayakody, and Jagdish Chandra each contained nontrivial prerecord elements, but the strength of the documentation, geographical proximity, and possible social pathways varied substantially.

The case of James Leininger, often rhetorically treated as decisive in popular discussion, was downgraded because of museum exposure, parental reinforcement, chronology issues, and the degree to which investigation and narrative development overlapped. Sujith Jayaratne was also downgraded because neighborhood links created ordinary information pathways. The lesson was not that the literature is worthless. It was that the evidential distribution is extremely uneven and that the strongest cases support anomalous information more directly than they support reincarnation specifically.

This is why the empirical tree gives the prerecord subset mild-to-moderate positive weight for continuity-friendly ontologies but refuses to treat the entire 2,500-case corpus as thousands of independent confirmations. The next genuinely useful empirical step would be a fixed forensic scorecard applied prospectively to the ten strongest prerecord cases, including leakage opportunities, chronology quality, miss rates, search-space size, and independence of documentation. The present revision has not inspected every primary case file or independently authenticated the prerecord chronology. Its small positive case weights remain conditional on that unresolved audit.

16.3 Simulation probability and the physics evidence

The total probability of simulation is constrained by the fact that quantum observer effects, Bell nonlocality, finite signal speed, Planck-scale structure, and emergent-spacetime ideas do not uniquely predict deliberate simulation. They weaken naive nineteenth-century material imagery more than they discriminate simulation from other nonfundamental accounts of spacetime. The analysis therefore gives the physics package only a modest simulation-specific likelihood advantage. Observer dependence is not evidence that one particular observer is privileged; Bell nonlocality is not a signature of a computer network; a finite signal speed need not be a processor clock; and Planck scales do not establish literal pixelation. The strongest inference from modern physics is that the manifest spacetime world may not be metaphysically fundamental, not that it was necessarily programmed.

Conversely, recent physical-limit arguments against full-universe simulation are narrower than their headlines. Bounds derived from the information and energy resources available under our physics constrain simulators implemented under physics like ours and simulating the entire universe at comparable resolution. They do not directly constrain every different-physics or selective-rendering architecture. Such alternatives still owe a resource budget for coherent experience, shared records, and reproducible experiments; moving to an interface does not make those costs vanish. These papers constrain particular full-resolution implementations. They neither disprove all generated simulation nor establish that participatory simulation is inexpensive.

16.4 The creator hypothesis and the problem of evil

Generic creator hypotheses and theistic versions of survival are not added as an independent headline branch because they overlap heavily with S1, idealism, and dualism. A benevolent creator raises survival conditional on that creator existing, but the creator hypothesis itself is affected by fine-tuning, hiddenness, religious diversity, suffering, lawfulness, and alternative structured explanations.

The problem of evil was especially important because its force is architecture-dependent. Under ordinary plural realism or plural idealism, the suffering of humans, animals, and prehistoric creatures represents an enormous quantity of actual phenomenal suffering. Under a one-subject world, most of that suffering is representational. This does not solve the problem of first-person suffering, and it does not show that a benevolent creator exists. It means only that the same apparent historical record has radically different moral likelihood under different assumptions about how many phenomenal subjects exist.

The paper therefore avoids the common mistake of treating the problem of evil as a single global likelihood against every consciousness-first or creator model. It is a conditional discriminator: severe first-person gratuitous suffering hurts benevolent single-subject hypotheses directly, while vast remote third-person suffering hurts benevolent plural-subject hypotheses much more strongly.

16.5 Current AI and what it does - and does not - show about consciousness

Contemporary AI makes several otherwise abstract distinctions vivid. Highly capable language and reasoning behavior can be generated by systems whose phenomenal status remains uncertain. This weakens any casual inference from intelligent behavior to consciousness and makes it easier to separate computation and cognition-like performance from phenomenality without assuming that current AI is definitely unconscious.

Current AI does not falsify functionalism. Serious functionalists do not claim that sheer complexity or benchmark performance automatically creates consciousness; they appeal to particular functional organizations. Nor does uncertainty about AI consciousness prove biological naturalism. The correct update is narrower: the gap between sophisticated function and confident evidence of phenomenality prevents behavioral success from serving as a shortcut to substrate independence.

The same distinction matters for the AI-operated S1 argument. An advanced AI need not itself be phenomenally conscious in order to operate a simulation. If functionalism is false, it could still be an extraordinarily powerful causal system that designs environments and connects genuine conscious participants to them. The data-generation argument therefore does not depend on attributing a human-like inner life to the simulator.

16.6 Local ontology versus ultimate ontology

A central conceptual simplification is the distinction between the ontology governing our local apparent world and the ontology of base reality. Discovering that our world is simulated would undermine the inference from local neural closure to ultimate physicalism even if base reality were entirely physical. Relative to the simulated world, a participant's mind can lie outside local physics; relative to base reality, the whole system may still be physical.

This is why the Bayesian tree classifies by the nearest death-relevant architecture. An S1 world is counted as S1 regardless of whether the external participant ultimately lives in a physical, idealist, or dualist base. Without that convention, the same world would be counted once under simulation and again under ultimate idealism or dualism. The local/ultimate distinction is therefore not merely philosophical bookkeeping; it is a dependence-control device.

Appendix A. Central Model Arithmetic

Start = (0.350, 0.250, 0.180, 0.220), in the order local production, participatory simulation, idealist interface, separable/transmission.

Brain factors = (1.00, 0.80, 0.68, 0.18).

Unnormalized structural masses = (0.350000000, 0.200000000, 0.122400000, 0.039600000); total = 0.712000000.

Structural posterior = (0.491573034, 0.280898876, 0.171910112, 0.055617978). Continuation conditionals = (0.02, 0.92, 0.68, 0.90). Structural continuation = 0.435213483.

Corroborated-perception row = (1.00, 1.20, 1.20, 1.20). Multiply this row with the other raw empirical rows in Section 4.7 before applying the discount.

Raw empirical products = (1.000000000, 1.215521236, 1.333323763, 1.078309575).

Final empirical multipliers = raw products raised to 0.75 = (1.000000000, 1.157635704, 1.240799799, 1.058175242).

Final unnormalized masses = (0.350000000, 0.231527141, 0.151873895, 0.041903740). Total = 0.775304776.

Final normalized architecture weights = (0.451435372, 0.298627260, 0.195889281, 0.054048086).

Final continuation = sum(normalized architecture weight × continuation conditional) = 0.465613776 = 46.5613776%, reported as about 47%. Arithmetic precision supports replication; it is not precision of belief.

Perception-row factor in each alternative branch	Continuation, decimal probability
1	0.437965588
1.1	0.452448938
1.2	0.465613776
1.3	0.477651459
1.4	0.488715725

The 1.10 comparison yields 0.452448938. Raising that benchmark to 1.20 adds 0.013164837, approximately one percentage point. The no-positive-row comparison yields 0.437965588; adding the 1.20 row adds 0.027648187, approximately three percentage points. These comparisons answer different questions and must not be conflated.

Retaining 75% of log evidence is equivalent to reducing its logarithm by 25%. The exponent is an explicitly chosen robustness device. It does not prove conditional independence or estimate the actual covariance of the underlying observations.

There is no added creator or solipsism probability mass. The companion metaphysics paper's 42.1% simulation scenario is not substituted into this separate prior. A joint reconciliation would require an explicit distribution over both papers' architecture variables, rather than averaging or copying their marginal percentages.

Sensitivity replication: participatory start $S = 0.025 + 0.75p$; production start $= 1 - S - I - D$. Multiply each start by its brain factor and final empirical multiplier, normalize, and weight by the continuation conditional. The mildly negative and stacked scenarios replace the complete final empirical multiplier; they do not apply an additional 0.90 row.

Scenario	Precise continuation
central	0.465613776
lower frequency	0.341225568
higher frequency	0.635674141
lower alternatives	0.362466863
combined	0.249172960
strict identity	0.400970313
negative	0.413709381
stacked negative	0.127956564
stacked positive	0.685499012

Appendix B. Observation Matrix

This matrix concerns compatibility and possible tests. Compatibility alone is not a likelihood advantage. A one-local-subject scenario also need not be philosophical solipsism about all reality: its simulator, external subjects, or wider world may still exist.

Observation	One local subject	Participatory simulation	Plural idealist interface
Lawful world and brain dependence	Requires persistent generative structure	Requires a specified interface	Requires psychophysical mapping
Autonomous other humans	Can be generated, with explanatory cost	Expected in multiplayer versions	Expected with many finite subjects
Advanced AI behavior	Relevant conditional on AI consciousness	Does not settle participant count	Does not settle subject count
Unobserved stable records	Requires storage beyond experience	Compatible	Compatible

Observation	One local subject	Participatory simulation	Plural idealist interface
Alleged shared anomalous information	Common generator can mimic sharing	Requires ordinary-channel exclusion	Requires ordinary-channel exclusion
Computational signature	Supports particular generated worlds if robust	May support simulation; does not identify participation	Not uniquely predicted
Apparent suffering of others	Less actual suffering only if others do not suffer	Depends on participants and operator	Problem for benevolent plural creation
Severe first-person suffering	Negative for benevolent generator	Negative for benevolent operator	Negative for benevolent generator

Reliable shared reports do not by themselves refute a sufficiently flexible one-subject generator. A test must discriminate specified models, not merely be compatible with more than one subject.

Appendix C. NDE Evidence Review and Documentation

The review consulted selected full-text sections of Long’s 2021 synthesis, the NDERF research index and individual accounts, Ring and Cooper’s original study, Cousins’s direct book review, Rush’s critique, Sartori’s response, Augustine’s methodological critique, and Martial et al.’s 2025 review. Long’s 2014 article, Long and Woollacott’s 2024 study, the original Sartori case, and AWARE studies were assessed through accessible abstracts or metadata together with the available discussions. Long’s commercial books were not read in full. The 2021 essay is a Bigelow competition submission, not a Missouri Medicine journal article. The review is targeted, not exhaustive.

Candidate selection prioritized specific potentially checkable observations, stronger corroboration, and distinct evidential categories. The following ten items are not ten verified paranormal cases or a certified ranking of every NDERF account. Several originated with independent researchers cited by Long.

Candidate	Evidence considered	Principal limitation
Sartori patient 10	Hospital observations, staff actions, oral cleaning; critique and response	Chronology and ordinary access disputed; not established irreversible death
AWARE 2014 awareness account	Prospective study with verifiable reported observations	Not a hidden visual-target success; exact brain-state timing uncertain
Nancy (Ring-Cooper)	Acquired blindness; independent witness interview and medical-history checks	Retrospective; auditory pathways and shared recollection remain possible
Vicki (Ring-Cooper)	Congenital blindness and reported visual experience	Central observations lacked locatable witnesses
Destiny C. (NDERF)	ICU resuscitation details; claimed next-morning written account	Clinical knowledge; original notes and independent staff records unavailable
Lori (NDERF)	Surgical complications allegedly known before explanation	Timing and corroboration self-reported

Candidate	Evidence considered	Principal limitation
Grady (NDERF)	Resuscitation and staff speech; claimed nurse confirmation	Heart-lung support complicates absent-perfusion claims; no independent record
Meg C. (NDERF)	Previously unknown deceased-relative information	No pre-confirmation record or independent family interview
Kate L. (NDERF)	Unusual surgical action and relative identification	Long reporting delay; contemporaneous corroboration unavailable
Karen T. (NDERF)	Reported observation and later recognition of unfamiliar surgeon	Retrospective identification; incidental prior information not excluded

Archive accounts are cited as testimony requiring verification. Their emotional force, number, and subjective certainty are not independent likelihood multipliers. Corroborated information could favor nonlocal information access without uniquely demonstrating a surviving personal subject. A suitable case audit records correct and incorrect details, chronology, witnesses, ordinary information pathways, and alternatives under each specified hypothesis.

Selected References

- Aru, J., Larkum, M. E., and Shine, J. M. (2023). The feasibility of artificial consciousness through the lens of neuroscience. *Trends in Neurosciences*, 46(12), 1008-1017.
- Augustine, K. (2022). How not to do survival research: Reflections on the Bigelow Institute essay competition. *Journal of Scientific Exploration*, 36(2), 366-398. <https://doi.org/10.31275/20222581>.
- Augustine, K., and Martin, M. (eds.) (2015). *The Myth of an Afterlife: The Case Against Life After Death*. Rowman and Littlefield.
- Bean, R. (2019). Cryptology from the crypt: how I cracked a 70-year-old coded message from beyond the grave. *The Conversation*, 29 August 2019.
- Bering, J. M. (2006). The folk psychology of souls. *Behavioral and Brain Sciences*, 29(5), 453-498.
- Block, N. (1978). Troubles with functionalism. *Minnesota Studies in the Philosophy of Science*, 9, 261-325.
- Bostrom, N. (2003). Are you living in a computer simulation? *Philosophical Quarterly*, 53(211), 243-255.
- Bostrom, N., & Shulman, C. (2023). *Propositions Concerning Digital Minds and Society*, version 1.21.
- Bourget, D., and Chalmers, D. J. (2023). Philosophers on philosophy: The 2020 PhilPapers Survey. *Philosophers' Imprint*, 23(11).
- Braude, S. E. (2003). *Immortal Remains: The Evidence for Life After Death*. Rowman and Littlefield.
- Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., Constant, A., Deane, G., Fleming, S. M., Frith, C., Ji, X., Kanai, R., Klein, C., Lindsay, G., Michel, M., Mudrik, L., Peters, M. A. K., Schwitzgebel, E., Simon, J., and VanRullen, R. (2023). Consciousness in artificial intelligence: Insights from the science of consciousness. *arXiv:2308.08708*.
- Carroll, S. M. (2011). Physics and the immortality of the soul. *Scientific American guest blog*, 23 May 2011.
- Carroll, S. M. (2016). *The Big Picture: On the Origins of Life, Meaning, and the Universe Itself*. Dutton.
- Carroll, S. M. (2022). The quantum field theory on which the everyday world supervenes. In O. Bueno, M. B. Fagan, and E. Castellani (eds.), *Levels of Reality in Science and Philosophy*, 27-46. Springer.

- Chalmers, D. J. (1995). Absent qualia, fading qualia, dancing qualia. In T. Metzinger (ed.), *Conscious Experience*, 309-328. Schöningh.
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Chalmers, D. J. (2022). *Reality+: Virtual Worlds and the Problems of Philosophy*. W. W. Norton.
- Chalmers, D. J. (2023). Could a large language model be conscious? *Boston Review*, 9 August 2023; arXiv:2303.07103.
- Churchland, P. M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, 78, 67-90.
- Cousins, W. E. (2012). Review of Evidence of the Afterlife: The Science of Near-Death Experiences, by Jeffrey Long with Paul Perry. *Journal of the Society for Psychical Research*, 76.1 (906), 53-57.
- Descartes, R., and Elisabeth of Bohemia (1643). Correspondence, May to June 1643. In L. Shapiro (ed. and trans.), *The Correspondence between Princess Elisabeth of Bohemia and René Descartes*. University of Chicago Press, 2007.
- Fischer, J. M., and Mitchell-Yellin, B. (2016). *Near-Death Experiences: Understanding Visions of the Afterlife*. Oxford University Press.
- Frankish, K. (2016). Illusionism as a theory of consciousness. *Journal of Consciousness Studies*, 23(11-12), 11-39.
- Gillogly, J. J., and Harnisch, L. (1996). Cryptograms from the crypt. *Cryptologia*, 20(4), 325-329.
- Gilmore-Bykovskiy, A., Kottler, H., Cochran, A., et al. (2025). A prospective observational study of lucid episodes in advanced dementia. *The Gerontologist*, 65(9), gnaf178. <https://doi.org/10.1093/geront/gnaf178>.
- Greyson, B., and van Lommel, P. (2023). AWAreness during RESuscitation and EEG activity. *Resuscitation*, 193, 110014 (letter), with reply by Parnia et al.
- Gurnee, W., Sofroniew, N., Pearce, A., et al. (2026). Verbalizable representations form a global workspace in language models. *Transformer Circuits*, July 6; arXiv:2607.15495, July 16. [Research report / preprint.] <https://transformer-circuits.pub/2026/workspace/index.html>.
- Hoffman, D. D. (2019). *The Case Against Reality: Why Evolution Hid the Truth from Our Eyes*. W. W. Norton.
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32(127), 127-136.
- James, W. (1898). *Human Immortality: Two Supposed Objections to the Doctrine*. Houghton Mifflin.
- Kastrup, B. (2019). *The Idea of the World: A Multi-Disciplinary Argument for the Mental Nature of Reality*. Iff Books.
- Kekecs, Z., Palfi, B., Szaszi, B., et al. (2023). Raising the value of research studies in psychological science by increasing the credibility of research reports: the Transparent Psi Project. *Royal Society Open Science*, 10(2), 191375. <https://doi.org/10.1098/rsos.191375>.
- Kelly, E. F., Kelly, E. W., Crabtree, A., Gauld, A., Grosso, M., and Greyson, B. (2007). *Irreducible Mind: Toward a Psychology for the 21st Century*. Rowman and Littlefield.
- Kipping, D. (2020). A Bayesian approach to the simulation argument. *Universe*, 6(8), 109.
- Leibniz, G. W. (1714). *Monadology*, section 17.
- Leslie, K., Chan, M. T. V., Myles, P. S., Forbes, A., and McCulloch, T. J. (2010). Posttraumatic stress disorder in aware patients from the B-Aware trial. *Anesthesia and Analgesia*, 110(3), 823-828.
- Linassi, F., Zanatta, P., Tellaroli, P., Ori, C., and Carron, M. (2018). Isolated forearm technique: a meta-analysis of connected consciousness during different general anaesthesia regimens. *British Journal of Anaesthesia*,

121(1), 198-209.

Long, J. (2014). Near-death experiences: Evidence for their reality. *Missouri Medicine*, 111(5), 372-380.
<https://pubmed.ncbi.nlm.nih.gov/25438351/>.

Long, J. (2021). Evidence for survival of consciousness in near-death experiences: Decades of science and new insights. Bigelow Institute for Consciousness Studies competition essay.
<https://www.nderf.org/NDERF/Research/EvidenceBigelow.pdf>.

Long, J., and Perry, P. (2010). *Evidence of the Afterlife: The Science of Near-Death Experiences*. HarperOne.

Long, J., and Woollacott, M. (2024). Long-term transformational effects of near-death experiences. *Explore*, 20(5), 103030. <https://doi.org/10.1016/j.explore.2024.103030>.

Martial, C., Cassol, H., Charland-Verville, V., Pallavicini, C., Sanz, C., Zamberlan, F., Vivot, R. M., Erowid, F., Erowid, E., Laureys, S., Greyson, B., and Tagliazucchi, E. (2019). Neurochemical models of near-death experiences: A large-scale study based on the semantic similarity of written reports. *Consciousness and Cognition*, 69, 52-69.

Martial, C., Fritz, P., Gosseries, O., Bonhomme, V., Kondziella, D., Nelson, K., and Lejeune, N. (2025). A neuroscientific model of near-death experiences. *Nature Reviews Neurology*, 21(6), 297-311.
<https://doi.org/10.1038/s41582-025-01072-z>.

Mashour, G. A., Frank, L., Batthyany, A., Kolanowski, A. M., Nahm, M., Schulman-Green, D., Greyson, B., Pakhomov, S., Karlawish, J., and Shah, R. C. (2019). Paradoxical lucidity: A potential paradigm shift for the neurobiology and treatment of severe dementias. *Alzheimer's and Dementia*, 15(8), 1107-1114.

McClelland, T. (2026). Agnosticism about artificial consciousness. *Mind & Language*, 41(4), 482-502.
<https://doi.org/10.1111/mila.70010>.

Moraes, L. J., Barbosa, G. S., Castro, J. P. G. B., Tucker, J. B., and Moreira-Almeida, A. (2022). Academic studies on claimed past-life memories: A scoping review. *Explore*, 18(3), 371-378.

Nahm, M., Greyson, B., Kelly, E. W., and Haraldsson, E. (2012). Terminal lucidity: A review and a case collection. *Archives of Gerontology and Geriatrics*, 55(1), 138-142.

Near-Death Experience Research Foundation (NDERF) (n.d.). Destiny C. experience account. Accessed October 6, 2026. https://www.nderf.org/Experiences/1destiny_c_nde.html.

Near-Death Experience Research Foundation (NDERF) (n.d.). Grady experience account. Accessed October 6, 2026. https://www.nderf.org/Experiences/1grady_nde.html.

Near-Death Experience Research Foundation (NDERF) (n.d.). Karen T. experience account. Accessed October 6, 2026. https://www.nderf.org/Experiences/1karen_t_nde.html.

Near-Death Experience Research Foundation (NDERF) (n.d.). Kate L. experience account. Accessed October 6, 2026. https://www.nderf.org/Experiences/1kate_l_nde.html.

Near-Death Experience Research Foundation (NDERF) (n.d.). Lori experience account. Accessed October 6, 2026. https://www.nderf.org/Experiences/1lori_nde.html.

Near-Death Experience Research Foundation (NDERF) (n.d.). Meg C. experience account. Accessed October 6, 2026. https://www.nderf.org/Experiences/1meg_c_nde.html.

Near-Death Experience Research Foundation (NDERF) (n.d.). Research overview and experience archive. Research index and the Destiny C., Lori, Grady, Meg C., Kate L., and Karen T. accounts, accessed October 6, 2026. https://www.nderf.org/NDERF/Research/Research_Overview_Right.htm.

Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.

Parnia, S., Keshavarz Shirazi, T., Patel, J., et al. (2023). AWAREness during REsuscitation II: A multi-center study of consciousness and awareness in cardiac arrest. *Resuscitation*, 191, 109903.

<https://doi.org/10.1016/j.resuscitation.2023.109903>.

Parnia, S., Spearpoint, K., de Vos, G., et al. (2014). AWARE: AWAREness during REsuscitation, a prospective study. *Resuscitation*, 85(12), 1799-1805.

Pinto, Y., Neville, D. A., Otten, M., Corballis, P. M., Lamme, V. A. F., de Haan, E. H. F., Foschi, N., and Fabri, M. (2017). Split brain: divided perception but undivided consciousness. *Brain*, 140(5), 1231- 1237.

Ring, K., and Cooper, S. (1997). Near-death and out-of-body experiences in the blind: A study of apparent eyeless vision. *Journal of Near-Death Studies*, 16(2), 101-147.
<https://digital.library.unt.edu/ark:/67531/metadc799333/>.

Robinson, W. (2023). Epiphenomenalism. *Stanford Encyclopedia of Philosophy*, substantive revision June 15, 2023; accessed September 30, 2026. <https://plato.stanford.edu/entries/epiphenomenalism/>.

Rush, M. J. (2013). Critique of A Prospectively Studied Near-Death Experience. *Journal of Near-Death Studies*, 32(1), 3-14.

Sanders, R. D., Gaskell, A., Raz, A., et al. (2017). Incidence of connected consciousness after tracheal intubation: A prospective, international, multicenter cohort study of the isolated forearm technique. *Anesthesiology*, 126(2), 214-222.

Sarraf, M., Woodley of Menie, M. A., and Tressoldi, P. (2021). Anomalous information reception by mediums: A meta-analysis of the scientific evidence. *Explore*, 17(5), 396-402. <https://doi.org/10.1016/j.explore.2020.04.002>.

Sartori, P. (2013). Response to critique of A Prospectively Studied Near-Death Experience. *Journal of Near-Death Studies*, 32(1), 15-36.

Sartori, P., Badham, P., and Fenwick, P. (2006). A prospectively studied near-death experience with corroborated out-of-body perceptions and unexplained healing. *Journal of Near-Death Studies*, 25(2), 69-84.
<https://digital.library.unt.edu/ark:/67531/metadc799351/>.

Schwitzgebel, E. (2008). The unreliability of naive introspection. *Philosophical Review*, 117(2), 245-273.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417-457.

Searle, J. R. (1992). *The Rediscovery of the Mind*. MIT Press.

Seth, A. K. (2025). Conscious artificial intelligence and biological naturalism. *Behavioral and Brain Sciences*, advance online article. <https://doi.org/10.1017/S0140525X25000032>.

Stevenson, I. (1968). The combination lock test for survival. *Journal of the American Society for Psychical Research*, 62(3), 246-254.

Stevenson, I. (1974). *Twenty Cases Suggestive of Reincarnation*, 2nd ed. University Press of Virginia.

Stevenson, I. (1997). *Reincarnation and Biology: A Contribution to the Etiology of Birthmarks and Birth Defects*, 2 vols. Praeger.

Stevenson, I. (2003). *European Cases of the Reincarnation Type*. McFarland.

Stevenson, I., Oram, A. T., and Markwick, B. (1989). Two tests of survival after death: Report on negative results. *Journal of the Society for Psychical Research*, 55, 329-336.

Sudduth, M. (2016). *A Philosophical Critique of Empirical Arguments for Postmortem Survival*. Palgrave Macmillan.

Thouless, R. H. (1948). A test of survival. *Proceedings of the Society for Psychical Research*, 48, 253-263.

Timmermann, C., Roseman, L., Williams, L., Erritzoe, D., Martial, C., Cassol, H., Laureys, S., Nutt, D., and Carhart-Harris, R. (2018). DMT models the near-death experience. *Frontiers in Psychology*, 9, 1424.

Tollock, M., Leontovich, N., Gonzalez, A., & Parnia, S. (2025). A multi-site prospective study of paradoxical lucidity in moderate to severe dementia. *Innovation in Aging*, 9(Supplement 2), igaf122.2914. [Conference abstract.] <https://doi.org/10.1093/geroni/igaf122.2914>.

Tucker, J. B. (2008). Ian Stevenson and cases of the reincarnation type. *Journal of Scientific Exploration*, 22(1), 36-43.

Tucker, J. B. (2016). The case of James Leininger: An American case of the reincarnation type. *Explore*, 12(3), 200-207.

van Lommel, P., van Wees, R., Meyers, V., and Elfferich, I. (2001). Near-death experience in survivors of cardiac arrest: a prospective study in the Netherlands. *The Lancet*, 358(9298), 2039- 2045.

Vazza, F. (2025). Astrophysical constraints on the simulation hypothesis for this Universe: why it is (nearly) impossible that we live in a simulation. *Frontiers in Physics*, 13, 1561873.

Vicente, R., Rizzuto, M., Sarica, C., Yamamoto, K., Sadr, M., Khajuria, T., Fatehi, M., Moien-Afshari, F., Haw, C. S., Llinas, R. R., Lozano, A. M., Neimat, J. S., and Zemmar, A. (2022). Enhanced interplay of neuronal coherence and coupling in the dying human brain. *Frontiers in Aging Neuroscience*, 14, 813531.

Walleczek, J., von Stillfried, N., Schmidt, S., Wittmann, M., Kirmse, K. A., Moll, J., & Kekecs, Z. (2025). Metascientific replication project with the advanced meta-experimental protocol of the transparent psi project procedures for testing the precognitive effect claimed by Bem. *PLOS ONE*, 20(11), e0335330. <https://doi.org/10.1371/journal.pone.0335330>.

Xu, G., Mihaylova, T., Li, D., Tian, F., Farrehi, P. M., Parent, J. M., Mashour, G. A., Wang, M. M., and Borjigin, J. (2023). Surge of neurophysiological coupling and connectivity of gamma oscillations in the dying human brain. *Proceedings of the National Academy of Sciences*, 120(19), e2216268120.

Appendix D. Generated-mind preservation and branching continuation

The distinction between participatory and generated simulation concerns how a subject is instantiated. It does not, by itself, determine what happens when the simulated life ends. The Bostrom-Yampolskiy interview is particularly useful because its discussion of uploading and preservation draws attention to causally connected continuation without assuming that strict numerical identity is settled. This appendix applies that distinction to the paper's probability model.

D.1 What preservation would have to accomplish

An operator might preserve a generated subject's relevant state and later resume the process, restore it in another environment, or sustain a conscious successor with sufficiently strong historical and psychological connections. Such a route can satisfy the paper's broad criterion even if a stricter identity theory rejects it. Merely keeping data, producing an unrelated look-alike, or claiming that information exists somewhere does not qualify.

Branching makes the distinction unavoidable. Two distinct successors cannot both be numerically identical to one original under ordinary identity logic. Both might nevertheless preserve psychologically important relations. The question of whether either is the original experiencing subject remains separate from whether either has memories, attachments, and projects continuous with the earlier life. A successor's sincere testimony is relevant to present experience and psychological fidelity; it cannot by itself decide the disputed identity theory.

Bostrom's practical attitude toward uploading is a reason to examine this broader survival criterion. It is not proof of subject transfer. Nondestructive copying also leaves a separate reason to protect the original: preserving a possible original subject is not made redundant merely by creating another morally important being.

D.2 Split the residual before changing its conditional

The benchmark groups 0.075 starting mass of generated simulation with 0.275 nonsimulation local-production mass. Both receive the local-production brain factor and empirical multiplier. This produces a final generated-simulation weight of $0.075 / 0.775304776 = 0.096736151$, approximately 9.67%.

Let g be the probability of qualifying continuation conditional on this generated-simulation branch and the evidence already represented. The benchmark sets $g = 0.02$, as it does for the rest of local production. Making g explicit, while leaving all architecture weights, empirical updates, and other continuation conditionals fixed, gives:

$$P(\text{continuation} \mid g) = 0.465613776 + 0.096736151 \times (g - 0.02).$$

This is a replacement of the generated branch's continuation conditional. It is not an extra independent probability mass, an additional NDE multiplier, or a second count of simulation.

The 2% allowance remains in nonsimulation local production.

Generated-mind continuation g	Overall continuation	Interpretation
0%	46.37%	No qualifying generated-mind continuation
2%	46.56%	Low-preservation benchmark
10%	47.34%	Modest preservation scenario
25%	48.79%	Intermediate scenario
50%	51.20%	One-half conditional scenario
75%	53.62%	Strong preservation scenario
100%	56.04%	Limiting case, not a forecast

The overall result crosses 50% when g exceeds approximately 37.55%. This threshold is a property of the stated model. It is not a measured probability of backups, a confidence interval, or a recommendation to assign that conditional. It shows why the below-half conclusion should be described as conditional on low generated-mind preservation, rather than robust to the mechanism now under discussion.

D.3 Ability, retention, choice, and continuity

The conditional g compresses several questions: does sufficient person-specific information survive; will someone instantiate a relevant process; will it be conscious; and will it satisfy the adopted continuity criterion? These must be represented jointly or by a chain of conditional probabilities. Multiplying separate unconditional guesses would incorrectly assume independence. Under a generated-consciousness hypothesis, the fact that the original simulation already hosted a subject also changes what should be assumed about an appropriate resumed implementation.

Governance matters. Ethical concern, a commitment to participants, scientific interest, or a desire to compensate suffering could motivate preservation. Resource constraints, abandonment, hostile goals, or indifference could defeat it. The interview supplies examples of morally intelligible motives; it does not tell us their prevalence among hypothetical operators. Preserving a life and providing a good subsequent life remain different outcomes, so welfare is not smuggled into the definition of continuation.

For the same reason, improvement in the case for computational consciousness need not always lower survival. At fixed total simulated mass, shifting weight from S1 with its 92%

continuation conditional to S2 with a lower g reduces continuation. But if g approaches or exceeds the S1 conditional, that comparison weakens or reverses. If the shift also changes total simulation prevalence, both effects must be calculated. The earlier direction-of-effect claim is therefore conditional, not an intrinsic disadvantage of generated minds.

D.4 Explicit comparison with the metaphysics estimate

After the benchmark's brain and empirical updates, total simulation mass is 39.54%, not its 32.5% starting allocation. The metaphysics paper's separate simulation scenario is 42.15%. To show a coherent bridge, impose 42.15% while keeping the participatory/generated ratio within simulation and the relative nonsimulation weights fixed. The resulting continuation values are 47.57% at $g = 2\%$, 49.95% at $g = 25\%$, and 52.52% at $g = 50\%$.

This proportional bridge is a sensitivity calculation. It does not establish a preferred joint prior or justify importing the metaphysics posterior as independent evidence. It exposes what must be chosen to reconcile the marginals, rather than averaging their headlines. The same evidence must not be applied again merely because it appears in a companion paper.

D.5 Revised interpretation and discriminating evidence

The benchmark remains reproducible at 46.56%; the expanded family no longer supports a robust statement that continuation is less likely than termination. No new preferred point estimate is adopted because neither the interview nor present evidence calibrates g . The paper now reports both the low-preservation benchmark and the explicit generated-preservation sensitivity.

Evidence capable of changing g would concern actual retention and restoration policies, fidelity of preserved states, conscious realization, and the relevant identity relation. Evidence that simulation is possible or that an advanced agent is benevolent would not settle all of those questions. Nor would a successful biological or digital revival automatically distinguish original-subject return from a conscious psychological continuer. Tests must specify which rival accounts predict different observations.

Sources and attribution

Source: Bostrom, Nick, interviewed by Roman Yampolskiy (2026). "Superintelligence Author: We May Already Have AGI | Nick Bostrom." The Roman Forum with Roman Yampolskiy, YouTube, October 16. <https://www.youtube.com/watch?v=m880ajt5sNM>. Interview discussion; conceptual claims and reported judgments, not an empirical calibration dataset.

Source: Parfit (1984), *Reasons and Persons*, for branching and psychological continuity; Bostrom and Shulman (2023), *Propositions Concerning Digital Minds and Society*, version 1.21, for digital survival and moral standing. <https://nickbostrom.com/propositions.pdf>. The five-way decomposition, sensitivity equation, threshold, and proportional bridge are this paper's analytical extensions. No interview timestamp or verbatim quotation is asserted.