

Cryonics and Personal Survival Across Metaphysical Frameworks

A Probabilistic Decision Analysis of Identity, Reconstruction, Afterlife Interaction, and Evidence-Conditioned Revival

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

With AI-assisted research and drafting; contributions and disclosures below

Circulation edition | September 24, 2026

RodMalpert.com

Abstract

Cryonics is often treated as a technological question: can a preserved brain someday be repaired? This paper asks the next question: if the repair works, has the person survived? Four issues must be kept separate. First, can the intended physical system be rebuilt and supported? Second, is the result conscious? Third, is the conscious subject numerically the same person? Fourth, does the result preserve enough of the original person's memories, attachments, projects, character, and agency to count as meaningful psychological continuation? The paper compares three revival pathways - repair of the original brain, repair that uses surviving information to reconstruct missing details, and digital uploading - across physicalism, idealism, dualism, and Russellian monism. In the reference model, the first two pathways do much better than conventional uploading, but the numbers are model outputs rather than measured survival rates. The paper then adds a complication that cryonics discussions often leave outside the model: what if some form of postmortem continuation already occurs? In that case revival could rescue a life that would otherwise end, add a return to embodied life, leave an independent continuation untouched, or in some scenarios displace a person from something better. That is why preservation and later revival are treated as separate decisions. The goal is not to prove cryonics works, prove an afterlife exists, or recommend enrollment. It is to make the assumptions visible enough that readers can substitute their own beliefs and values.

Reading guide

Sections 1-3 introduce the problem in ordinary language: what cryonics is trying to preserve, the main revival pathways, and the difference between identity and psychological continuation. Sections 4-7 ask what changes under four broad metaphysical frameworks. Section 8 turns those assumptions into a transparent numerical model. Sections 9-11 ask how the analysis changes if postmortem continuation is possible and if later patients can learn from earlier revivals. Section 12 considers welfare and identity after revival. Readers interested mainly in the argument can skim the equations and use the surrounding prose and tables; the appendices preserve the full technical specification.

Interpretation of the numbers: the percentages in this paper are not observed frequencies and are not actuarial forecasts for any cryonics patient. They show what follows from a declared set of philosophical and probabilistic assumptions. The value of the model is that those assumptions can be inspected and changed. The paper also does not evaluate any individual's preservation quality, contract, or medical choices.

1. What the paper is actually asking

Cryonics is the low-temperature preservation of a person after legal death in the hope that future technology may be able to repair the damage and restore function. No cryonics organization can currently revive a cryopreserved patient (Alcor 2026). This paper deliberately sets the engineering question to one side. It asks a conditional philosophical question: suppose future technology can rebuild a functioning system from a preserved patient. What, exactly, would have to be true for that success to count as the person's survival?

The paper keeps four questions separate. First: was the intended physical system successfully reconstructed and made functional? Second: is the result conscious? Third: is that conscious subject the very same person, rather than merely a new being with the same memories? Fourth: does the result preserve enough of the original person's

memories, character, attachments, projects, and agency to count as meaningful psychological continuation? The technical appendix abbreviates these four ideas as T, C, I, and M, but the main text will normally use the words. A technology could succeed physically while still failing at consciousness, same-person identity, or psychological continuation. Keeping those questions separate prevents the word 'revival' from doing too much work.

A simple example shows why the distinction matters. Imagine a future system that speaks exactly as the patient would and remembers the patient's life. Physical reconstruction may have succeeded. But if the system has no experience at all, consciousness has not returned. If it is conscious but is a newly created copy, same-person identity may have failed. Yet even a new subject might preserve a great deal of what the patient cared about, which is why psychological continuation is tracked separately. The paper does not assume in advance that strict identity or psychological continuation is the only thing that matters.

The same standard for meaningful psychological continuation is used under every metaphysical theory. A psychologically rich continuer counts whether physicalism or idealism is true. Conversely, a theory could allow the original soul or subject to persist while losing nearly all of the psychology that made the earlier life recognizable. In that case same-person identity could hold while psychological continuation fails. They are therefore different questions (Parfit 1984; Lewis 1976; DeGrazia 2005).

For calculation, the model treats meaningful psychological continuation as a yes-or-no threshold, even though real continuity comes in degrees. The threshold means substantial, person-specific continuity, not perfect recall of every event. There is no validated neuroscience cutoff at which a percentage of preserved structure becomes a percentage of preserved personhood. Section 12 therefore returns to the more realistic graded questions of connectedness, agency, and welfare.

1.1 Scope and how to read the evidence

The baseline case is a patient whose brain has not already suffered massive destruction before preservation. That does not mean the brain is pristine: loss of circulation, delay, disease, and the preservation process can all cause damage. Cases involving severe dementia, extensive head injury, or extremely poor preservation may present much larger information losses and are not modeled here.

Different sources do different jobs in this paper. Neuroscience and cryobiology can tell us something about what biological structures and functions can survive severe interruption. Philosophy supplies competing accounts of consciousness and personal identity. Decision theory helps compare choices under uncertainty. None of those sources, by itself, tells us the probability that a particular cryonics patient will survive.

The model therefore uses explicit judgment calls for ontology weights, local mechanisms, and revival outcomes. Those numbers are not survey results and were not taken from the philosophers cited. Agreement among AI systems is also not independent empirical calibration. The Bayesian structure is useful because it forces the assumptions into the open and lets readers change them; it should not be mistaken for a medically calibrated posterior probability of survival.

2. Three ways revival might happen

The paper compares three analytically different routes to revival. They are not predictions about which technology will arrive first. The distinction matters because each route asks a different philosophical question.

Table 1. Revival pathways and the separate embodiment axis.

Case	What is physically restored	What still has to be shown
Repair the original brain	The preserved biological brain remains the cognitive organ and recoverable damage is repaired.	Does consciousness resume, and is the conscious person still the original with meaningful psychological continuity?
Repair plus anchored reconstruction	The original brain remains the anchor, but missing details are reconstructed from surviving person-specific evidence.	Is the surviving anchor strong enough for identity and meaningful psychological continuation?
Digital upload	A scan-derived implementation runs on conventional digital hardware.	Is the implementation conscious, and if so does it preserve the original person or only a copy/continuer?

Case	What is physically restored	What still has to be shown
Embodiment (separate issue)	A biological body, robotic body, or virtual environment can be paired with any pathway.	How well does the embodiment support adaptation, communication, agency, and welfare?

The important difference between original-brain repair and anchored reconstruction is not how many atoms or cells get replaced. A large repair could be tightly constrained by surviving information, while a small injury could destroy something uniquely important. Anchored reconstruction therefore concerns how much of the restored person must be inferred rather than directly recovered. Gradual replacement of functioning neurons is a different thought experiment: it tests continuity across an ongoing process rather than reconstruction after missing information. Chalmers's distinctions among gradual, destructive, nondestructive, and reconstructive uploading help keep these cases apart (Chalmers 2010, sections 9-11; 2014).

2.1 A brain-computer interface is not an upload

A repaired biological brain could live in a repaired body, control a robotic body, or receive its sensory world through a computer interface. In all three cases, the biological brain can remain the cognitive system. Changing the brain's inputs and outputs is not the same thing as copying the person into software. This matters because retaining the original brain may preserve forms of biological and causal continuity that a conventional upload does not.

A new body or environment could still change the revived person's experience. Hormones, bodily sensations, sensory systems, and patterns of action all help shape psychology. Those effects matter to the quality of continuation, but they do not automatically mean that a different person now exists. The paper therefore treats embodiment as a separate issue from the three revival pathways.

A newly built biological or neuromorphic brain is a boundary case. If the original brain can be repaired, there is no need to duplicate it merely to connect it to a new environment. If the original cannot function but still contains useful information, a future system might reconstruct a successor in biological or brain-like hardware. Calling hardware 'neuromorphic' does not, by itself, show that it has the concrete properties required for consciousness.

2.2 Damage, missing information, and reconstruction

Three questions are easy to blur together: Is the tissue damaged? Does the surviving evidence still tell us what the damaged part used to be like? And does the missing detail matter for personal continuation? A future repair may have to infer some missing information without simply inventing a new person. Conversely, a system can behave fluently even when some of its autobiographical content is wrong.

Merkle's information-theoretic approach focuses on recoverability: important structure may remain inferable even when present medicine cannot restore it (Merkle 1992). That is a criterion for thinking about information loss, not proof that a future technology could reconstruct an individual mind. A damaged book gives the basic intuition. A torn page may be repairable; a missing sentence may have one strongly supported reconstruction; or several different sentences may fit equally well. Those are importantly different cases even if every finished version reads smoothly.

What matters is not the disappearance of ongoing electrical activity as though the person were stored in a current that simply drains away. The harder question is which structural and dynamic features must survive, or remain inferable, for the relevant cognitive organization to restart. A wiring diagram alone may not be enough; synaptic strengths, cellular states, molecular information, or other dynamics may matter. Current preservation research does not provide a validated formula converting 'percent structure preserved' into 'percent memory or personality preserved' (McKenzie et al. 2024).

2.3 What current biology shows - and does not show

Some major interruptions of brain activity are reversible. During deep hypothermic circulatory arrest, measured electrocerebral silence can be followed by recovery of electrophysiological activity (Stecker et al. 2001a, 2001b). That is more relevant than ordinary sleep as an example of interrupted function. But the comparison has strict limits: an isoelectric scalp EEG does not prove that every brain process has stopped, and a brief surgical interruption is very different from decades of cryogenic storage.

Recent cryobiology is also encouraging in a narrower sense. German and colleagues (2026) reported recovery of selected hippocampal functions after vitrification of adult mouse tissue, including tissue prepared from whole brains after rewarming. The study measured properties such as cellular excitability, synaptic transmission, and long-term potentiation. It did not revive a whole animal, demonstrate a conscious functioning whole brain, or recover known autobiographical memories. The authors also reported important limitations, including low-yield preparations and possible selection toward responsive samples (German et al. 2026; Maya-Romero and Zoncu 2026).

Vitrification reduces ice-crystal damage, but it does not eliminate every source of injury. Cryoprotectant toxicity, osmotic stress, uneven perfusion, and thermal damage remain possible. Preserving structure is also different from showing that the structure can currently be returned to normal function. Reviews such as McKenzie et al. (2024) make the case that preserved information may have future value while acknowledging that reconstruction remains speculative.

The fair conclusion is modest. Biology gives reasons not to equate temporary loss of function with destruction of the person-relevant information. It also gives reasons to think structural preservation can improve. What it does not yet give us is an evidence-based estimate of how much person-specific inference a future revival would require, or whether such revival will become possible at all.

3. What counts as the same person?

3.1 Copies, branching, and the identity problem

If one person is copied into two distinct conscious successors, ordinary numerical identity says that both cannot literally be the one original person. But that logical point does not answer the practical question of whether one or both successors preserve relationships that mattered about survival.

Parfit made psychological connectedness and continuity central to this distinction, while Cerullo develops an explicitly permissive account of branching identity (Parfit 1984, Part III; Cerullo 2015). Other theories resist that move. Nozick's closest-continuer view, non-branching psychological theories, and Lewis's treatment of fission preserve a larger role for uniqueness or identity (Nozick 1981; Lewis 1976; Shoemaker 1984). There is therefore no neutral philosophical rule that makes the copying problem disappear.

The practical position here is narrower than any full theory of identity. If the original remains viable, creating a conscious copy is not automatically a justification for destroying the original. A reader who values psychological continuation may regard a copy as a meaningful survivor; a reader who requires the continued original organism or subject may not. Either way, a conscious successor has moral interests of its own.

3.2 The continuity-through-inference argument

The continuity-through-inference argument developed in this project begins with an ordinary fact: people change. Education, grief, conversion, illness, relationships, and decades of experience can substantially alter personality and priorities without making us say that the earlier person died. More than one future development could have grown out of the same earlier life, and many of those futures would still count as that person's continuation.

That suggests a limited defense of inference-assisted repair. A reconstruction does not have to reproduce one perfectly determined psychological snapshot in every detail. If several reconstructions are strongly constrained by the surviving person-specific evidence, each might preserve enough memory, commitments, relationships, capacities, and practical orientation to count as meaningful continuation. Rapid technologically assisted change should not be treated as death merely because ordinary change would have occurred more slowly.

The argument is conditional. If an ordinary psychological change of a certain size would preserve meaningful psychological continuation, and if a reconstruction preserves comparably strong person-specific connections without adding some further relevant break in causal or phenomenal continuity, then the reconstruction should not be rejected merely because it was fast or partly inferential. That does not answer a theory that independently requires the continued existence of a particular organism or nonphysical subject.

3.3 Where the analogy breaks

Reconstructing the past is not exactly like living into the future. In ordinary life, the new experience really happens to the person. A reconstruction could instead install an apparent memory of something that never occurred. A plausible personality assembled from population averages would not be a reconstruction of this individual. The inference therefore needs a substantial causal anchor in the preserved person, and recovered facts, supported inferences, and genuine uncertainties should remain distinguishable.

Causal history matters for the same reason. A randomly generated duplicate could resemble someone perfectly without having grown out of that person's actual life. Repair constrained by the original brain and reliable personal records has a historical connection even when some details are uncertain. Philosophers can still disagree about whether that connection is enough for strict identity, but it is a strong candidate for something that matters to psychological continuation.

There is no defensible rule that says five percent inferred psychology is harmless while fifty percent is fatal. We do not even have a unit in which such percentages could be measured. The importance and pattern of what is preserved matter more than a raw count. Losing a few central relationships or memories could matter more than changing many peripheral habits.

A useful future assessment would therefore ask three questions: How much person-specific causal and informational structure remains? How wide is the range of reconstructions compatible with that evidence? And how much meaningful psychological connection would each candidate preserve? Multiple plausible candidates are not automatically disqualifying. They become troubling when they diverge sharply on features central to the person's life. Where the evidence is genuinely silent, preserving the uncertainty may be better than manufacturing a confident false memory.

3.4 Metaphysics changes what repair must accomplish

The physical amount of damage does not become larger or smaller when we change metaphysical theories. What changes is what the repaired system must do in order to count as the person's continuation. A reductionist physicalist may accept extensive but well-anchored reconstruction. A biologically conservative physicalist may require more organismic continuity. A theory with a separable subject may allow strict identity to survive large psychological changes but still require that the same subject reconnect with the repaired body.

The model therefore separates three uncertainties that are often collapsed: whether consciousness resumes, whether the same subject resumes, and how much of the person's psychology is preserved. Its anchored reconstruction reference case assumes a strong surviving anchor. Section 8.4 then shows what happens as that assumption is weakened. The model tries not to charge the same underlying uncertainty three times under different labels.

4. If physicalism is true

Physicalism says, roughly, that a person's mental life ultimately depends on physical reality rather than on a separately existing immaterial soul. That still leaves major disagreements about what makes a person persist. A physicalist may emphasize the continuing brain, the whole living organism, psychological continuity, functional organization, or phenomenal continuity (Unger 1990; Olson 1997; Dainton 2008).

For a brain-continuity physicalist, repairing and restarting the original biological brain is the most straightforward cryonics case. There is no separate immaterial subject that must be found and brought back. But even here, restoring generic biological function is not enough. A living but radically altered brain could be conscious while failing to preserve the memories, attachments, and psychology that make the revival valuable.

anchored reconstruction - repair that fills gaps from surviving evidence - is easiest for psychological or reductionist forms of physicalism to accept. Material replacement is not automatically fatal if the right person-specific structure and causal history remain. A theory that demands an uninterrupted conscious process can object, but it then has to explain why ordinary periods in which consciousness is not being exercised do not also break identity.

4.1 Physicalism does not guarantee digital consciousness

Some physicalists think the right functional organization is enough for consciousness, regardless of substrate. Others think biological or other concrete dynamical properties are essential. Seth's biological naturalism is a contemporary example of skepticism that computation alone is sufficient, while still allowing artificial consciousness to become more plausible as systems become more brain-like or life-like (Seth 2025).

That means uploading has two separate hurdles. First: is the digital system conscious at all, rather than merely a highly accurate model of a person? Second: if it is conscious, is it the same person or a psychologically meaningful continuer? The first question cannot be settled by saying that a simulation is 'only a simulation,' because functionalists deny that the relevant computation merely depicts the mental process; they argue that it implements it.

Survey data illustrate how unsettled these questions are, but they do not determine this paper's probabilities. In the 2020 PhilPapers survey, respondents divided sharply over psychological, biological, and further-fact accounts of identity and over whether mind uploading would count as survival (Bourget and Chalmers 2023). Those survey answers come from different respondent groups and are not conditioned on this paper's assumptions, so they are evidence of disagreement rather than inputs to the numerical model.

Recent AI evidence modestly strengthens the case that some of the organization associated with conscious access is substrate-general. Anthropic's 2026 J-space work reports a global-workspace-like set of internal representations in Claude: reportable, controllable, reasoning-relevant, capacity-limited, and broadly broadcast. That is relevant to the possibility that a future digital reconstruction could support access-like cognition. It is not, however, evidence that phenomenal consciousness or personal identity is thereby preserved.

The S1/S2 distinction from the companion metaphysics work clarifies the identity issue. An S2-like reconstruction can generate a conscious digital subject without showing that the preserved patient's original subject continued; an S1-like account would instead require a persisting subject outside the reconstructed local process. J-space bears primarily on the first question - whether digital organization can reproduce consciousness-relevant function - and only indirectly on the second. The reference model is therefore not numerically recalibrated on J-space alone.

4.2 Animalism and the whole-body question

Animalism identifies the person with the human animal rather than with a detachable psychology. On that view, preserving the organism may matter in a way that preserving the brain alone does not. Whole-body preservation could therefore retain a philosophical route that neuropreservation forecloses. But merely keeping more tissue does not prove that the same organism survives a long period of preservation and reconstruction (Olson 1997; van Inwagen 1990; Parfit 2012).

There is therefore no universal philosophical premium for whole-body preservation. It depends on how much probability one gives to organism-based identity and on whether the relevant preservation procedure actually maintains the continuity that theory requires. More tissue is not automatically better if preserving it comes at the expense of the information most important to brain function and personal continuity.

5. If idealism is true

Idealism makes mind or consciousness fundamental rather than derivative from matter. That does not mean every individual personality automatically survives death. For cryonics, the key question is how a finite person is individuated. On some idealist views, restoring the right organization may lawfully reconstitute the person. On others, an independently persisting subject would have to reconnect with the restored organism. The paper treats those as different mechanisms rather than pretending that the word 'idealism' settles the issue.

5.1 Reconstitution in idealist theories

Kastrup's dissociative idealism treats living organisms as the outward appearance of dissociated mental processes within a wider consciousness (Kastrup 2017, 2023). He argues that ordinary silicon computation is unlikely, by itself, to produce private subjectivity and points to metabolism as the known outward correlate of dissociation. He does not claim to have logically ruled out artificial consciousness.

The cryonics application is therefore conditional. If restoring a biological organism also restores the relevant dissociative process, renewed subjectivity may follow without a technician having to retrieve a separate soul. But one universal consciousness does not make every finite person identical. We still need an account of why this particular finite perspective, with this history, returns rather than merely some conscious process.

A reconstituted subject might preserve meaningful psychological continuation even if strict finite-person identity remains disputed. A theory that requires the very same finite subject needs an additional rule of individuation. Cases in which dormant dissociative identities later reappear in living people are suggestive analogies at most; they do not establish what happens after severe biological deterioration and reconstruction.

5.2 Conscious agents and artificial interfaces

Hoffman's conscious-agent framework also starts with consciousness as fundamental and treats spacetime as an interface rather than ultimate reality (Hoffman and Prakash 2014; Prentner and Hoffman 2024). That makes artificial interfaces conceptually possible, but it does not answer whether ordinary silicon systems are conscious or whether the original subject of a preserved person would be the one connected to a future implementation.

Work on fusions of conscious agents likewise warns against two easy simplifications: the theory is not just physicalist combination in different vocabulary, and it is not simply one undifferentiated cosmic mind in which personal identity is automatically solved (Hoffman, Prakash, and Prentner 2023). A cryonics application still needs a bridge from the repaired interface to the particular person's experiential history.

5.3 The receiver-and-broadcaster problem

A radio analogy captures the hardest problem for theories that require reattachment. Perfectly repairing a receiver does not prove that the original broadcaster still exists, is still transmitting, or will reconnect to that receiver. Likewise, even a perfectly restored brain might not guarantee the return of a separately persisting subject. The model therefore treats 'reattachment' as its own uncertain event rather than hiding it inside technical success.

Several very different reattachment rules are possible. Return might occur automatically by law, because the restored structure matches the old subject, because the subject chooses to return, or because some external operator decides. If the mechanism is unspecified, neither near-zero nor near-certain return is justified. Section 9 shows why the difference matters for consent and for possible displacement from a postmortem state.

Three outcomes are then possible: the original subject returns; a different conscious subject occupies the reconstructed psychology; or no conscious subject is associated with the system. The first can satisfy same-person identity. Either of the first two can satisfy meaningful psychological continuation if the psychological and causal links are strong enough. A reader who cares only about strict original-subject return should therefore focus on the same-person identity column rather than redefining meaningful psychological continuation.

5.4 What cryonics adds under idealism

If the finite subject would otherwise cease, preserving the structures needed to reconstitute that subject could be crucial. If the subject continues independently, cryonics may add something different: a return to this biography, this embodiment, and these relationships. And if the independent continuation is much better than revived life, return could be a loss. None of those conclusions follows from idealism alone. Section 9 models the continuation and welfare questions separately.

6. If dualism is true

Dualism covers several importantly different views. Some posit a distinct subject or soul that can exist apart from the body. Others hold that consciousness has irreducible properties while remaining tightly law-governed by brain organization. Hylomorphism, emergent dualism, interactionism, and epiphenomenalism should not be treated as one doctrine. The model therefore asks what local mechanism each version would require for revival.

6.1 Separable subjects and hylomorphism

On a separable-subject view, the brain is not the ultimate subject. Preserving it may preserve the conditions through which the subject expressed itself, but successful repair still leaves a further question: does the same subject

reconnect? Swinburne's account of the soul and personal identity is a prominent example of this family of views (Swinburne 1997). It does not, by itself, assign a probability to cryonic reconnection.

Hylomorphic theories tie the soul to the form of a living organism. That can make organismic continuity important without yielding a simple answer about a cryopreserved adult, an isolated head, or a later reconstructed body. The narrower point needed here is simply that some theories make preservation of the organism itself part of what survival requires.

Ordinary resuscitation shows that temporary loss of observable function does not always mean permanent personal departure. It does not show that a much longer interruption preserves the same subject-body relation. Likewise, defining death as irreversible departure of the soul cannot prove that a future reconstruction retrieves that soul. The missing bridge is exactly what a cryonics theory would need to explain.

6.2 Emergent selves and property dualism

Hasker's emergent dualism treats a distinct conscious self as arising from an organized physical system (Hasker 1999). Rebuilding a suitable system could therefore allow a conscious subject to emerge without anyone 'manufacturing a soul.' But whether it would be the earlier subject is a separate question. Restarting the generator is an analogy, not an identity proof.

Property dualism can instead hold that irreducible phenomenal properties arise according to psychophysical laws. Chalmers's organizational-invariance arguments offer one route by which consciousness might be preserved across different substrates without reducing consciousness to physics (Chalmers 1996, chapter 7). Other property dualists can reject that route. Strict epiphenomenalism also shows why irreducible consciousness does not automatically imply survival after brain function ends.

For that reason, original-brain repair and strongly anchored reconstruction remain comparatively favorable under brain-linked forms of dualism. Uploading depends on what kinds of organization the psychophysical laws recognize. Separable-subject theories add the further reconnection problem. The particular numbers in Section 8 are illustrative model choices, not consequences of the label 'dualism.'

7. If Russellian or neutral monism is true

Russellian monism starts from the idea that physics describes the structure and behavior of reality while leaving open something about the intrinsic nature of what has that structure. That intrinsic nature may matter to consciousness. Nearby views include panpsychism, panprotopsyism, cosmopsychism, and forms of neutral monism, but they disagree sharply about how individual subjects arise (Chalmers 2016; Goff 2017; Shani 2015).

Neutral monism is related but not identical: it often treats both mental and physical descriptions as arising from a more basic neutral reality. The companion metaphysics paper groups these positions pragmatically for modeling while acknowledging the boundaries among them (Malpert 2026a). The present paper keeps that approximately 40 percent family weight so the cryonics analysis remains comparable to the companion work.

7.1 Organization and intrinsic nature

Three questions matter here. Is abstract functional organization enough for consciousness? Does consciousness require a more specific concrete causal organization? And whatever creates consciousness, what makes one conscious subject the same subject over time? A theory can answer the first two while leaving the identity question open.

The fact that carbon brains and silicon computers are made from the same basic physical universe does not show that they instantiate the same consciousness-relevant organization. Chemical, electromagnetic, and cellular relations may matter. At the other extreme, ordinary material turnover makes a demand for the exact same particles implausible. What may matter is a type of material, causal organization, or continuing pattern rather than particle identity itself.

Integrated information theory provides a useful comparison because it emphasizes intrinsic causal structure rather than abstract software alone (Tononi and Koch 2015). It is not itself a Russellian theory, and it does not prove that digital devices are unconscious. The point is methodological: on any structure-sensitive view, we have to ask what the hardware actually does, not merely whether the same program description can be written for it.

7.2 Local and global individuation

Some theories build a subject largely from local relations among parts; others begin from a larger conscious whole and explain how a finite perspective becomes differentiated within it. The first family resembles brain-linked revival. The second resembles a consciousness-first or reconstitution model. Neither automatically settles personal identity.

A locally constituted subject might cease when its organizing relation breaks. A global theory might instead contain a lawful rule by which the same finite perspective can recur. Both can support meaningful psychological continuation if the revived consciousness is richly connected to the earlier person's psychology. Both still need an additional account before consciousness alone establishes same-person identity.

Original-brain repair has an advantage on these views because it preserves the original organized system rather than asking a new implementation to realize the relevant relations. Anchored reconstruction depends on how much of those relations survive or can be reliably reconstructed. Digital uploading depends on whether the actual digital hardware realizes them or merely models them. Under concrete-structure assumptions the repair/upload gap can be large; under permissive functional assumptions it can nearly disappear.

8. Putting the possibilities into one model

Nontechnical reading option. The rest of Section 8 is the bookkeeping layer. Readers can skip the equations and still follow the argument by reading Tables 2-4 and the surrounding prose. From this point onward, the main text uses plain-English labels wherever possible. Compact letters and abbreviations are retained mainly inside equations and in Appendix A so that the arithmetic remains reproducible.

8.1 Two layers: ultimate ontology and revival mechanism

The companion metaphysics paper supplies four working weights: physicalism about 26 percent, Russellian monism about 40 percent, idealism about 23 percent, and dualism about 10 percent (Malpert 2026a). Because those rounded numbers sum to 99 rather than 100, the model normalizes them. These are inherited judgment weights, not measured frequencies of metaphysical truth, and the categories are forced to be exclusive for bookkeeping even though philosophical usage sometimes overlaps.

The model then adds a second layer: an external-participant architecture in which a pre-existing subject participates in the world through an avatar or interface. This possibility can sit on top of more than one ultimate ontology, so it is not an extra metaphysical category to be added on top of the first four. Bostrom's familiar simulation argument mostly concerns generated simulated observers; an externally participating subject is a different architecture, closer to Chalmers's distinction between different kinds of simulated beings (Bostrom 2003; Chalmers 2022). In the equations and appendix this architecture is abbreviated S1.

Outside the external-participant case, the model groups revival behavior into three mechanisms. In a brain-linked mechanism, consciousness is realized through the brain or brain-like organization. In holistic reconstitution, a finite perspective is reconstituted from a larger or holistic mind. In an enduring-subject mechanism, a persisting subject exists separately and must reattach. Appendix A abbreviates these as BL, HR, and ES. Those letters are useful for the formal specification, but the main discussion will use the descriptive names.

$$J(o, S1) = w(o)s(o); \quad J(o, k) = w(o)[1 - s(o)]m(o, k), \quad k \text{ in } \{BL, HR, ES\}.$$

The equation above is only bookkeeping: it makes the ontology layer and the local-mechanism layer add to one without counting the same possibility twice. The symbols correspond to the plain-English categories just described, and Appendix A lists the actual inputs. If a future simulation or external operator used a different kind of hardware, backup system, or restoration rule, the response probabilities would have to change rather than being hidden in a new unnormalized branch.

The reference model assigns the external-participant architecture a 25 percent share within each ontology. This is a simplifying scenario, not a result imported from the companion papers. The metaphysics simulation scenario is about 42%, covering both generated and participatory minds. The survival benchmark implies about 39.5% total simulation, including about 29.9% participatory mass. These differently conditioned quantities are compared explicitly in Appendix C rather than treated as identical.

8.2 Comparing pathways, not predicting current patients

For readers who want the formal expression, the equation below computes the standardized score for either same-person identity or meaningful psychological continuation under each revival pathway. Readers who do not need the algebra can treat it as a weighted average over the worlds already described.

$$Q_j(Y) = \sum_o w(o) \{ [1 - s(o)] t_{oj}(Y) + s(o) u_j(Y) \}.$$

The crucial limitation is that these are comparisons among successfully constructed future systems. They are not estimates of the chance that a person entering cryonics today will reach that stage. Preservation quality, technological feasibility, institutional survival, selection for revival, and implementation failure are outside these standardized scores.

If successful revival were actually observed in the future, that evidence could also change our beliefs about which metaphysical or local mechanism is true. A full Bayesian update would then use the likelihood of that observation under each theory. The present table does something simpler: it holds the starting worldview weights fixed so the philosophical consequences of the three pathways can be compared cleanly.

8.3 From consciousness to identity and psychological continuation

Within brain-linked and holistic-reconstitution cases, the model first asks whether the restored system is conscious and then assigns separate conditional chances to same-person identity and meaningful psychological continuation. Within enduring-subject and external-participant cases, it distinguishes three possibilities: the original subject returns, a different conscious subject appears, or no subject appears. The first can satisfy same-person identity. Either conscious subject can satisfy meaningful psychological continuation if the required psychological connections are strong enough.

For simplicity, the reference model initially gives the same psychological-fidelity value to an original returning subject and to another conscious continuer. That is plainly contestable if an intervening afterlife transforms the original subject. Section 9.8 relaxes the simplification. Again, same-person identity and psychological continuation are not the same outcome: the same subject can return with little psychological continuity, or a different subject can preserve much of the earlier psychology.

8.4 What the reference model says

Table 2. Reference outcomes. Each cell is: same-person identity % / meaningful psychological continuation %.

Framework or mixture	Original-brain repair	Anchored reconstruction	Digital upload
Physicalism	84.2 / 96.0	79.2 / 91.1	22.0 / 53.4
Russellian monism	82.6 / 94.3	77.8 / 89.4	14.0 / 34.0
Idealism	63.9 / 75.7	61.2 / 71.8	13.6 / 24.3
Dualism	70.4 / 82.8	67.2 / 78.5	25.7 / 54.6
Combined, including external-participant architecture	69.3 / 81.5	66.2 / 77.3	24.1 / 43.7

Framework rows exclude the external-participant architecture; the combined row includes it. Read each pair as 'identity first, psychology second.' One decimal place is retained for reproduction, not to imply empirical precision.

The headline result is easy to state. Conditional on successful construction, original-brain repair scores about 69 percent for same-person identity and 81 percent for meaningful psychological continuation. Strongly anchored reconstruction scores about 66 and 77 percent. Conventional digital uploading scores about 24 and 44 percent. The corresponding consciousness estimates are about 84 percent for the two repair pathways and 45 percent for digital uploading. These are outputs of the reference assumptions, not observed survival rates.

The model can also ask how often same-person identity and meaningful psychological continuation hold together. Because both require consciousness, their joint probability is bounded even without assuming that identity loss and psychological loss are independent. For original-brain repair the joint value lies between about 67 and 69 percent; for anchored reconstruction between 59 and 66 percent; and for digital uploading between 23 and 24 percent. Those bounds are mathematical consequences of the chosen marginals, not new empirical evidence.

Table 3. Declared sensitivity cases. Each cell is: same-person identity % / meaningful psychological continuation %.

Scenario	Original-brain repair	Anchored reconstruction	Digital upload
Reference	69.3 / 81.5	66.2 / 77.3	24.1 / 43.7
Low attachment	57.2 / 68.1	54.1 / 64.5	13.9 / 32.1
High attachment	81.4 / 93.2	78.3 / 88.4	38.2 / 57.3
Shared identity minus 0.10	63.0 / 81.5	59.8 / 77.3	21.5 / 43.7
Weak reconstruction anchor	69.3 / 81.5	44.0 / 46.2	24.1 / 43.7
Reconstruction anchor as strong as original-brain case	69.3 / 81.5	69.3 / 81.5	24.1 / 43.7
Attachment-heavy idealism	60.2 / 71.8	57.4 / 68.1	23.0 / 42.3
Upload-compatible realization and identity	69.3 / 81.5	66.2 / 77.3	69.3 / 81.5
No external-participant architecture	77.4 / 89.3	73.2 / 84.7	17.2 / 38.9
External-participant share at 10%	74.2 / 86.1	70.4 / 81.7	20.0 / 40.8
External-participant share at 30%	67.7 / 79.9	64.8 / 75.8	25.5 / 44.7
Digital-upload return reduced to enduring-subject level	69.3 / 81.5	66.2 / 77.3	19.1 / 38.9

These are changes to assumptions, not confidence intervals. Appendix A defines the settings. Different-sized interventions should not be ranked as if they were a formal variance decomposition.

Sensitivity analysis matters more than the second decimal place. Changing how readily a persisting subject reattaches, how much weight is given to the external-participant architecture, or how strongly anchored reconstruction remains tied to surviving person-specific information can move the results substantially. The size of a sensitivity effect depends on how wide a range we choose to test, so the table should not be read as proving a single 'largest driver.'

One assumption is especially visible in uploading. Nearly half of the digital-upload identity score in the reference case comes from allowing an external participant to return to a digital continuer just as readily as to the original biological brain. An external simulator might impose substrate-specific rules instead. Equal return across substrates is therefore a testable assumption inside the model, not something that follows from simulation theory.

The small gap between original-brain repair and anchored reconstruction is also conditional. The reference reconstruction case assumes very strong surviving person-specific information: an 80 percent conditional same-person-identity value and a 92 percent conditional psychological-continuation value. Table 4 is therefore more informative than the single reference row because it shows how anchored reconstruction changes as the surviving anchor becomes weaker or stronger.

Table 4. Anchored-reconstruction sweep. Each cell is: same-person identity % / meaningful psychological continuation %.

Anchor index	0	0.25	0.50	0.75	1.00
Anchored reconstruction	44.0 / 46.2	50.4 / 55.0	56.7 / 63.8	63.0 / 72.7	69.3 / 81.5

The index moves linearly from a weak anchor to an anchor as strong as the original-brain reference case. At the low end, the conditional identity/psychology values are 45/55 percent; at the high end they are 85/97 percent. The reference values lie close to the high end. This index is not a percentage of neurons or memories retained.

If the external-participant architecture remains 25 percent overall but is allowed to fall unevenly across the four ontologies, original-brain repair still stays in a fairly narrow range in this particular audit. At the other extreme, if one assumes digital hardware is as consciousness-friendly and identity-friendly as the original-brain pathway, digital uploading can be made to match original-brain repair. The model therefore supports a narrower claim: source-brain repair is a broad hedge across the reference assumptions, not a theorem that every coherent upload theory must do worse.

8.5 How the companion afterlife estimate fits

The companion life-after-death paper reports about 47 percent broad continuation in its low-preservation benchmark (Malpert 2026b). This is neither a probability of heaven nor a strict-identity estimate. Its generated-mind extension permits higher or lower results under explicit assumptions. Appendix C distinguishes that model from the conditional revival scores and no-preservation counterfactual used here.

This model defines B0 as broad continuation without cryonic preservation. Its mechanism-weighted reference total is 43.87 percent. That is a separate counterfactual scenario, not the companion paper's current 46.56 percent result and not independent confirmation of it. The gap reflects different mechanism mixtures and conditioning. Substituting the companion headline would not supply the joint distribution needed to recalculate rescue and interference; Appendix C provides the bridge and limits.

It would be wrong to subtract the companion survival percentage from the revival score and call the remainder the value of cryonics. The percentages describe different events. What matters is their overlap in the same worlds: when revival adds continuation, when it adds embodiment to continuation already occurring, and when it could interfere with something better.

9. What if there is an afterlife?

The easiest way to read this section is to remember one question: compared with what? If death would otherwise end the person, revival can add continuation. If the person would continue in some other form, revival may add an embodied return, do nothing to the other state, coexist with it, or interfere with it. The model separates those cases rather than assuming that either 'afterlife' or 'revival' has one fixed effect.

A possibility of postmortem continuation changes the cryonics question in a surprisingly important way. The issue is no longer just 'Can revival work?' We must also ask what would have happened without revival, whether that state is good or bad, whether the person is actually experiencing it when revival occurs, whether revival changes it, and whether the same subject is involved. Those are separate questions.

The analysis therefore compares preservation and no-preservation for the same hypothetical person in the same underlying world. It allows many outcomes: useful revival, failed revival, two continuers that coexist, displacement from a postmortem state, later restoration, or a person who changes during the interval. None of the transition probabilities in this part of the model are established by empirical evidence; they are explicit scenario assumptions.

9.1 Three stages: the human life, a possible postmortem state, and revival

Table 5. Three stages and three distinct relations.

Relation	Question	Role
Human life -> postmortem state	Does the earlier human meaningfully continue in a postmortem subject or state?	The companion paper's broad-continuation target.
Postmortem state -> revived person	Does that subject return, re-individuate, coexist with a continuer, or remain separate?	Interaction between an existing continuation and revival.
Human life -> revived person	Is the revived subject the same person, or at least a rich psychological continuer?	The two central personal-survival targets of cryonics.

There are three stages to keep separate: the person during ordinary human life; any intervening postmortem subject or state; and the revived being. The relation between each pair can differ. The human and postmortem subject might be numerically the same person, while the postmortem subject changes so much that the revived person is psychologically closer to the earlier human than to that later state. Or the revived being might preserve the human's memories while the postmortem subject continues elsewhere. Same-person identity and psychological connectedness therefore cannot be collapsed into one line.

The 'postmortem state' is only an analytic label. The paper does not assume that an afterlife runs on the same clock as cryogenic storage. Thirty years in storage might correspond to no experienced interval, an immense interval, intermittent experience, or something to which ordinary duration does not apply. Any calculation about delay or opportunity cost has to specify that mapping rather than assume it.

9.2 Different afterlife models interact differently with revival

Table 6. Revival does not have the same meaning in every continuation model.

Mechanism	Without preservation	Possible effect of revival
Brain-linked	Usually no qualifying continuation; the reference leaves a small residual.	Repair may create continuation; external backups would change the baseline.
Holistic reconstitution	A life's experiential history may remain accessible within a larger mind.	Re-individuation may preserve, partition, or disrupt access; not necessarily literal soul recall.
Enduring subject	A separable subject may persist independently.	The original may return, another subject may preserve the earlier psychology, or neither may occur.
External participant	The participant may persist after avatar death.	The interface may reopen or not; refusal, reassignment, parallel continuation, or disruption depend on the rules.

A holistic-reconstitution model can create costs even if 'recalling a soul' is the wrong description. Re-forming a limited individual perspective might remove access to a richer global state. Conversely, return under an enduring-subject or external-participant model could be beneficial if the subject freely chooses it, keeps access to the better state, or gains something important by returning. The mechanism alone does not determine whether revival helps or harms.

9.3 What the reference calculation actually shows

The model first calculates a 'rescue' contribution: the chance that cryonics produces meaningful psychological continuation in worlds where the broad postmortem continuation counted by the companion paper would otherwise not occur. The word rescue is only shorthand. A world outside that broad-continuation category is not necessarily absolute nothingness; it simply lacks the particular kind of continuation being counted here.

It then calculates 'candidate exposure': the chance that broad postmortem continuation would already occur and that the original subject is also returned in either the enduring-subject or external-participant branches. Candidate exposure is not the same thing as being forcibly pulled out of heaven. For actual displacement, the postmortem subject must still be active when revival happens, the same subject must return, and the return must interfere with the better state. The reference calculation deliberately stops earlier and should therefore be read as an exposure measure, not a harm rate.

In plain English, actual displacement requires all of the following: an independent postmortem continuation exists; it is still active when revival occurs; the same subject returns; and that return interferes with the earlier state. The formal expression for those events is retained in the technical appendix.

The central calculation makes one additional simplifying assumption: within each model cell, whether broad postmortem continuation occurs is treated as independent of the revival outcome. That assumption is used for the point estimate, not asserted as a fact about reality. Section 9.4 removes it and asks how far the results can move while the same marginal probabilities are preserved.

Table 7. Where the reference rescue and possible-afterlife-interference numbers come from.

Mechanism	Weight	Continuation without cryonics	Psychological continuity after revival	Rescue contribution	Possible afterlife-interference exposure
Brain-linked	48.9%	2%	96.0%	46.0 points	0.0 points
Holistic reconstitution	16.5%	68%	87.3%	4.6 points	0.0 points*
Enduring subject	9.6%	90%	58.2%	0.6 points	3.9 points
External participant	25.0%	92%	58.2%	1.2 points	10.4 points
Total	100%	43.9%	81.5%	52.3 points	14.2 points

How to read Table 7: 'Model weight' is the share of the reference model assigned to each mechanism. 'Continuation without cryonics' is the chance that meaningful postmortem continuation already occurs in that mechanism. 'Psychological continuity after revival' is the chance that original-brain repair produces a meaningful continuer. The last two columns show how many percentage points each

mechanism contributes to the overall rescue total and to the narrower possible-afterlife-interference exposure. The 52.3% and 14.2% totals are not overall probabilities that cryonics is beneficial or harmful.

Under the reference assumptions, original-brain repair contributes 52.32 percentage points of rescue and 14.25 percentage points of candidate exposure to possible afterlife interference. In worlds without the paper's broad postmortem continuation, original-brain repair preserves meaningful psychology about 93 percent of the time and same-person identity about 81 percent. In worlds where broad continuation already occurs, the corresponding values are about 66 and 54 percent. Those are properties of the chosen mixture, not evidence that every no-afterlife theory favors cryonics or every afterlife theory opposes it.

The simple ratio of rescue contribution to candidate exposure is 3.67 for original-brain repair, 3.48 for anchored reconstruction, and 1.99 for digital uploading. If the model also counts possible interference with holistic reconstitution, the original-brain exposure measure rises and the ratio falls to about 2.15. That is useful because it shows how much the conclusion depends on what kinds of interaction are counted; it does not prove that every holistic revival is harmful.

9.4 Letting dependence vary

The point estimate above pairs the no-preservation and revival outcomes in one particular way. But different correlations are possible. The formal bound below asks a simple question: if we keep the same overall probability of broad postmortem continuation and the same overall probability of meaningful psychological continuation after revival, what is the smallest and largest possible overlap between 'no broad continuation' and 'successful psychological continuation'? Readers can skip the formula and use Table 8.

$$\max(0, m - b) \leq P(\text{not } B_0 \text{ and } M_j) \leq \min(m, 1 - b).$$

The program performs the same exercise for the other relevant events while enforcing the logical rule that same-person identity and meaningful psychological continuation both require consciousness. This yields coherent best- and worst-case couplings rather than changing unrelated totals independently.

Table 8. Original-brain-repair sensitivity to how the counterfactual worlds are paired.

Quantity	Reference pairing	Range if dependence varies
Rescue contribution (%)	52.3	49.1 to 55.2
Possible afterlife-interference exposure: enduring subject + external participant (%)	14.2	12.6 to 15.6
Rescue / exposure ratio	3.67	3.15 to 4.37
Possible exposure if holistic reconstitution is also counted (%)	24.4	22.2 to 26.8
Rescue / broader-exposure ratio	2.15	1.83 to 2.49

These bounds vary only the dependence between 'what would happen without cryonics' and 'what happens after revival.' They keep the worldview weights, mechanism shares, and revival-response probabilities fixed. They are therefore a narrow robustness check, not an overall uncertainty interval.

Within this limited dependence audit, the rescue contribution remains larger than the candidate-exposure contribution. But that does not settle the overall decision. The model remains much more uncertain about the weight of the external-participant architecture, whether the original subject is available to return, how good or bad the postmortem state is, how consent works, and how future institutions behave.

9.5 When could a better afterlife outweigh revival?

Probability alone does not answer the decision. Suppose a desired revival is valuable in a world where broad postmortem continuation would not occur, but suppose an existing postmortem state is much better than revived life in worlds where it would occur. Then even a relatively small chance of harmful displacement could matter if the value difference is large. The simplified calculation below isolates that tradeoff.

Ignoring other benefits, ordinary bad revivals, and costs only for this illustration, the expected difference is:

$$\Delta U = S V_R - E h(V_A - V_R).$$

$$h^* = (S/E) / (V_A/V_R - 1), \text{ for } E > 0 \text{ and } V_A > V_R.$$

This equation is intentionally incomplete. It ignores ordinary bad revivals, financial and institutional costs, partial harms, and many benefits. Its purpose is to show the break-even logic transparently: the more valuable the postmortem state is relative to revived life, the smaller the harmful-displacement probability needed to offset the rescue benefit. The letters in the equation are only compact labels for those quantities.

Table 9. Break-even harmful share of the candidate-exposure cases under the restricted original-brain example.

Value of postmortem state / value of revived life	2	5	10	100	1,000
Break-even harmful share of candidate exposure	No crossing	0.918	0.408	0.0371	0.0037

When the postmortem state is only twice as valuable as revived life, the raw break-even calculation is above 100 percent, so there is no crossing within the feasible range. A program must not clip that result to 100 percent and call it break-even. Costs or ordinary bad-revival losses can still reverse the decision even when this two-term slice has no crossing.

In this restricted example, if the postmortem state is only twice as valuable as revived life, even a 100 percent harmful share of the candidate-exposure cases is not enough to reverse the sign. If the postmortem state is ten times as valuable, the harmful share needed for break-even falls to about 41 percent. If it is one hundred times as valuable, the threshold falls to about 3.7 percent. Those are not universal thresholds. Adding holistic interference, ordinary revival harms, or costs can lower them further.

A reader who cares about strict numerical identity rather than psychological continuation can substitute the identity probability into the same calculation. That changes the rescue term and the threshold; identity is not a second benefit that can simply be added on top of psychological continuation.

The reader-substitution model goes one step further. It asks the reader for three inputs: the reader's own probability of a good afterlife that revival could displace; the chance of a desired revival if that afterlife is absent; and the effective probability that the good afterlife would be fully displaced. The compact equation below labels those quantities with letters, but Table 10 states the result in words. The point is to let readers see how their own beliefs change the decision without pretending that everyone shares the paper's reference values.

$$d^* = (1 - a)p_{V_R} / [a(V_A - V_R)].$$

Table 10. Reader substitution: how afterlife belief changes the break-even displacement threshold.

Reader belief in a good afterlife	Break-even total displacement	Break-even harmful share if return occurs 45% of the time	Interpretation
10%	80.0%	1.78; no crossing at $r = 0.45$	An 80% overall displacement probability would reach break-even.
45%	10.86%	0.241	A moderate displacement probability can matter.
90%	0.99%	0.022	A roughly 1% overall displacement probability offsets the hedge.
95%	0.47%	0.010	Less than one-half percent overall displacement reaches break-even.

The third column asks: if the original subject returns 45 percent of the time, what fraction of those returns would have to be harmful to reach break-even? A value above 100 percent means there is no crossing under that specified return probability, not immunity to every possible form of displacement.

A reader who assigns 95 percent to an excellent heaven can therefore test that belief directly without accepting the paper's 43.87 percent broad-continuation model. The calculation does not assume all Christians share the same credence or that heaven permits displacement. A reader who expects a neutral or bad postmortem state should use different values and may regard revival as helpful even when postmortem continuation occurs. Belief in continuation, the kind of continuation, and its value are three different inputs.

9.6 Who controls return?

If the postmortem subject can make an informed, competent, uncoerced choice about returning - and if refusal leaves that subject safely in the postmortem state - then consent is a major protection. A subject who prefers that state can simply decline. But this is a conditional safeguard, not a proof that voluntary return is always harmless. The choice could be misinformed, manipulated, impaired, or burdened by costs that occur before consent is possible.

If an external operator controls return, the operator's knowledge, motives, and powers matter. A specifically benevolent operator who blocks harmful returns could make displacement risk very low. Generic theism or generic simulation does not by itself imply such a rule. Automatic law and structural matching are additional possibilities. The model therefore treats the return mechanism as something to specify, not something to smuggle in under the word 'afterlife.'

The return rate and the value of the postmortem state may also be related. If the postmortem state is better, a genuinely free subject might be less willing to return. Holding the return rate fixed while making the postmortem state better is useful as a sensitivity test, but it should not be mistaken for the prediction of a voluntary-choice model.

9.7 Memory, exit, and temporary displacement

Remembering a postmortem state could help a revived person recognize that return was unwanted. But it does not follow that ending the revived life would restore the earlier state. Any 'exit and return' scenario requires additional assumptions about agency, medical decisions, reliable information, and the metaphysical rules governing what happens next. This paper does not recommend death as a route to an afterlife.

Memory is also not the only way a revived person could learn what happened. External evidence might inform someone who remembers nothing, while vivid memories could be false or impossible to act on. Even a reversible return could impose distress, lost time, or other opportunity costs. The model should therefore reduce harm only by the fraction that a real, ethically available reversal would actually avoid.

The more general expression below formalizes that point. It allows a return to impose both a long-run loss that may be partly reversible and a transition cost that cannot be undone. The simple h term used earlier is just the special case in which the loss is all-or-nothing.

A further possibility is an afterlife-memory 'firewall': experiences in a postmortem state occur but become inaccessible when the subject again operates through an ordinary human brain. Some metaphysical views make such a rule conceivable, but none establishes its probability. A revived person could sincerely remember nothing and still have displaced something valuable. Conversely, neuroscience showing that the person sincerely lacks a memory would not prove that the external afterlife event occurred.

9.8 What if the person changes in the afterlife?

Suppose the postmortem state includes a vast or radically expanded life. The returning subject could be numerically identical to the earlier human while becoming psychologically remote from that earlier self. Or the revived person could be reset to something much closer to the earlier human while losing much of what the postmortem subject had become. Neither outcome is captured well by a single identity label. The result is a three-stage problem of continuity and consent.

A person's advance instructions during ordinary human life therefore do not automatically exhaust the interests of a later postmortem or revived subject. If the same persisting subject can communicate a current, informed refusal, that refusal matters morally. When communication is impossible, future trustees face genuine uncertainty and should not assume their own preferences are identical to the subject's.

The reference psychological-continuation values assume ordinary, well-anchored return. A serious transformation scenario should replace them with branch-specific values: one for the original subject after a postmortem state, another for any other continuer, and further distinctions for whether that state occurred and whether its memories remain accessible. That preserves the meaning of psychological continuation instead of pretending the earlier 81.5 percent already covers every imaginable afterlife transformation.

10. Preserve now, decide later

Cryonic preservation and later activation are not the same decision. If preservation keeps the option open, future evidence can change whether revival should be attempted. That is a standard value-of-information point under uncertainty and irreversibility (Howard 1966; Miller 1975; Arrow and Fisher 1974). The option has value only if the future institution can actually wait, learn, and act responsibly.

10.1 Can early revivals teach later patients?

Yes, potentially - but not automatically. If an early revived person gives a credible account of a valuable postmortem state and describes revival as an unwanted loss, later patients might rationally delay activation. One report would not settle the question: reconstruction error, confabulation, timing of memory formation, and individual variation are alternative explanations. Independent corroboration or successful preregistered predictions would be much stronger evidence.

Future neurotechnology may help distinguish deliberate fabrication, memory integrity, and present experience. It cannot by itself prove that a sincere memory corresponds to an external afterlife rather than a hallucination, reconstruction artifact, or shared source of error. Evidence should therefore be separated into at least four layers: what the subject experiences, whether the report is sincere, whether it contains independently verifiable information, and which metaphysical explanation best accounts for that information.

10.2 Why repeated 'no afterlife memory' reports may not settle the issue

A person who remembers no afterlife tells us less than it first appears. Perhaps no afterlife occurred. Perhaps the original subject did not return. Or perhaps an afterlife occurred but the relevant memories cannot pass through the revived brain. To show the difference, the model compares a shared 'memory firewall' that blocks all relevant memories across cases with ordinary independent forgetting in each case. The compact equation below gives the technical learning curve; readers can skip it and go directly to Tables 11 and 12.

$$L_n = \phi + (1 - \phi)(1 - \rho)^n; \quad P(G \mid n \text{ nulls}) = a L_n / [1 - a + a L_n].$$

If there is a nonzero chance of one shared firewall applying to everyone, repeated null reports eventually stop lowering the afterlife probability very much. The posterior approaches a floor because the same single explanation can account for all of the failures. That floor belongs only to this model of observation; a different experiment could still be informative.

Table 11. After many null reports, the remaining afterlife probability if there might be one shared memory firewall.

Prior probability of afterlife	Shared firewall chance 10%	Shared firewall chance 50%	Shared firewall chance 90%
0.10	0.011	0.053	0.091
0.45	0.076	0.290	0.424
0.90	0.474	0.818	0.890

If blocking instead happens independently for each person, repeated null reports become much more powerful. Ten people who all independently fail to remember an afterlife are harder to explain by repeated bad luck than ten people whose memories are all blocked by the same universal rule. Table 12 shows how dramatically the two learning curves can diverge even when the first case looks identical.

Table 12. Identical single-case sensitivity, different long-run learning.

Number of null reports	Afterlife probability: shared firewall	Afterlife probability: independent blocking
0	0.4500	0.4500
1	0.3293	0.3293
3	0.2920	0.1502
10	0.2903	0.0049

Both examples start with a 45 percent afterlife probability. They assume an 80 percent chance that a person would report the afterlife if memory were available and, in the shared-firewall example, a 50 percent chance that one universal memory block applies. These are likelihood assumptions, not estimates of actual afterlife memory.

In the example, ten null reports leave the shared-firewall model near a 29 percent posterior, while the independent-blocking model falls below one-half percent. Neither number says which model nature actually follows. The point is that evidence from early revivals cannot be interpreted without a theory of how likely afterlife information is to survive revival in the first place.

10.3 Who volunteers first?

If everyone waits for somebody else to generate evidence, nobody goes first. That resembles a volunteer's dilemma (Diekmann 1985). But low belief in an afterlife does not make someone an ethically appropriate experimental subject. A person can be unconcerned about displacement and still face serious technical, welfare, identity, or institutional risks.

Early participation should therefore be genuinely voluntary and supported by favorable preservation, strong safeguards, clear preferences, and a willingness to contribute evidence. Independent review, patient representation, communication support, welfare monitoring, and stopping rules matter. Religious or metaphysical belief should not substitute for informed consent or clinical suitability.

Early volunteers may also be unrepresentative. Better-preserved or more recently preserved patients could be chosen first, and future procedures could differ across cohorts. Evidence from the first revived patients therefore has to be transported cautiously to later patients rather than mechanically applied to everyone.

10.4 Evidence-threshold instructions and institutional control

A preference not to be among the first experimental revivals is more useful when stated as an evidence threshold than as a numbered place in a queue. An instruction might ask future custodians to wait for independently reviewed evidence about earlier outcomes, memory reports, welfare, communication, and the ability of revived people to refuse further intervention. Success with a materially different procedure would not automatically satisfy the threshold.

Current institutional documents also matter because a personal preference is useful only if someone has authority and incentives to honor it. A public Alcor form posted in 2024 gives the organization substantial good-faith judgment about revival and rehabilitation and incorporates amendable general terms. That does not establish what any particular member's executed agreement legally requires. It shows why the actual contract, amendment rules, funding, representation, and successor authority need separate review (Alcor 2024, sections 1.1, 4.4, 5.2).

The proposed instruction is therefore a policy idea, not a legal opinion. Whether it can be made binding would require review by qualified counsel and the cryonics organization. The paper does not assume that a future custodian would automatically be bound by a preference written today.

10.5 The value of information - and the cost of waiting

A simple example shows why waiting can matter. Suppose the world is either a harmful-displacement state A or a favorable-revival state N. If acting immediately has slightly negative expected value, a later signal can still make activation attractive in the favorable branch and unattractive in the harmful branch. The numbers below are intentionally hypothetical; they illustrate the logic of learning, not the diagnostic accuracy of future testimony.

In the worked example, a positive signal makes the harmful state overwhelmingly likely and revival unattractive; a negative signal makes the favorable state much more likely and revival attractive. Being able to condition the decision on that signal creates positive expected value. This is the familiar value of sample information: useful information lets later decision-makers choose different actions in different states.

Real evidence is not free. Early subjects can be harmed, experiments consume resources, selected cases may be unrepresentative, and institutions can create pressure to act. Those costs belong in the calculation rather than being treated as though information simply appears.

Waiting is not free either. Institutions can fail, preserved material can be lost, governance can deteriorate, or a future opportunity can disappear. A sensible sequential policy compares activation now, waiting, and never activating. It is not a rule to postpone forever regardless of the evidence.

11. From probabilities to a decision

11.1 Comparing policies, not slogans

The decision variable is not simply 'cryonics: yes or no.' It is a policy: preserve or do not preserve; if preserved, under what evidence threshold and with what rules for activation, welfare, and consent? The expected-value

expression below compares the full trajectory under a preservation policy with the full trajectory without preservation, state by state.

$$\Delta U(\pi) = \sum_{\theta} P(\theta) \{E[U \mid \text{preserve}, \pi, \theta] - E[U \mid \text{no preservation}, \theta]\} - \text{cost}.$$

Where independent postmortem continuation is possible, the no-preservation baseline is not automatically zero. Where preservation itself could delay or alter that continuation, even 'preserve but never activate' may differ from ordinary death. If cryonics creates a parallel conscious successor while the original continues elsewhere, whether that successor's life counts as a benefit to the original person is a separate normative question.

The rescue/exposure calculations are only one slice of this larger decision. They leave out ordinary bad revival, preservation and custody costs, benefits inside continuation worlds, the chance that future institutions never reach activation, and many other losses. A reader who requires both same-person identity and meaningful psychological continuation should use their joint probability or bounds rather than adding the two percentages as if they were independent rewards.

11.2 Bad revival is not a new objection

Earlier cryonics literature already recognizes that revival can go badly. Moen's finite-upside argument treats cryonics as potentially rational when the alternative is terminal, but also discusses identity and unfavorable revival conditions (Moen 2015). Story argues against a simple 'nothing to lose' dominance claim because revival could be inescapably bad or radically transformative (Story 2024; see also Shaw 2009; Minerva 2018; Paul 2014).

The numerical example makes one basic point. If successful technology sometimes produces a very good revived life and sometimes produces a severely bad conscious life, the expected value depends on both probabilities and the severity of the bad outcome. A 50 percent chance of technical success does not determine whether cryonics is attractive. With the example's mild adverse outcome the expected gain is positive; making that adverse outcome ten times worse turns it negative.

Other decision rules can also matter. A risk-averse person may care more about severe downside than expected value alone captures. Rights-based constraints may rule out some experiments even when expected value is positive. Ambiguity aversion may penalize poorly understood risks. The paper therefore does not treat one favorable expected-value calculation as a recommendation to enroll or to cancel.

11.3 A long revived life versus delayed postmortem continuation

A revived life lasting centuries could be enormously valuable. But if revival delays access to an even better postmortem state, longer duration can also increase the opportunity cost. A fair comparison has to specify the same time horizon, quality of experience over time, discounting if any, whether the two states can coexist, and what happens after the revived life eventually ends.

The simple value ratios in Section 9.5 work only when the value assigned to the postmortem state and the value assigned to revived life are finite summaries of comparable alternatives. There is no general rule that adding more revived years scales benefits and losses in the same way. A timeless afterlife or literally infinite utility requires a different framework; infinity cannot be dropped into a finite break-even table and expected to behave normally.

Amnesia does not erase opportunity cost merely because the revived person cannot remember what was lost. At the same time, the actual interests and preferences of a contented revived person cannot simply be ignored because someone hypothesizes an unseen better state. The purpose of the model is to expose that conflict, not to resolve it by assuming that death automatically returns the person to the earlier state.

12. What happens after revival?

12.1 Two kinds of identity drift

Pre-revival drift concerns how the person may change between ordinary human life, a possible postmortem state, and revival. Post-revival drift concerns what happens after revival as enhancement technologies change memory, cognition, embodiment, or social integration. In both directions, strict identity can remain intact while direct psychological connectedness becomes weaker (Parfit 1984; Schechtman 1996; DeGrazia 2005).

Extreme mental speed is mainly a concern for uploads and heavily augmented systems, not an ordinary repaired biological brain. Faster thought or better memory is not inherently identity-destroying; it could strengthen continuity. The real concerns are unwanted rewriting, loss of valued relationships or commitments, forced merger, and changes so rapid that the person cannot understand or endorse them.

12.2 Revival first, enhancement later

Restoration, orientation, communication, enhancement, acceleration, and merger should be treated as separate decisions. The first goal after revival should be to establish comprehension, communication, welfare, and agency - not to maximize performance immediately. Radical changes should require separate information and consent, with reversible trials where genuinely possible.

The enhancement literature and group-mind scenarios give reasons to take those transitions seriously (Agar 2014; Sotala and Valpola 2012). A 2026 qualitative BCI study by Gilbert, Burkhart, and Morrill describes changes in experienced ownership, agency, and integration among implant users. The authors do not claim to have demonstrated literal ontological reconstitution. The study is relevant to lived self-conception, not a measured probability that a person ceases to exist.

12.3 Verifying continuity and protecting uncertain subjects

A future system saying "I am the original" would not prove numerical identity: the statement could arise from inherited memories. Source-specific tests can still assess psychological fidelity. A preservation archive should distinguish verified records, withheld test information, and uncertain reconstructed details. Installing an answer and then counting its repetition as independent confirmation would be circular. Appendix C separates the evidential targets of early revival tests.

Uncertainty about strict identity does not erase the moral standing of a conscious successor. A being with the patient's memories, attachments, hopes, and fears has interests even if philosophers dispute whether it is numerically the same person. The same precaution applies during incomplete restoration if suffering is possible before reliable communication. Long and Resnik (2025) argue for precaution in post-mortem brain research. That supports welfare oversight without assuming that every unresponsive preparation is conscious.

13. Conclusions

Cryonics is not one problem. It is at least five: physical reconstruction, consciousness, strict identity, meaningful psychological continuation, and the decision whether revival improves the person's total outcome. Keeping those questions separate makes the debate clearer. Repairing the original brain avoids several extra assumptions required by conventional uploading, while strongly anchored reconstruction can come surprisingly close if enough person-specific information survives.

Under the reference assumptions, original-brain repair yields 69.3 percent for same-person identity and 81.5 percent for meaningful psychological continuation; anchored reconstruction yields 66.2 and 77.3 percent; and conventional digital uploading yields 24.1 and 43.7 percent. Those figures are scenario outputs, not empirical patient-survival probabilities. Their value is diagnostic: they show which assumptions create the gap. If digital hardware is assumed to support consciousness and identity as readily as the original brain, the gap can disappear; if the surviving person-specific anchor is weak, reconstruction falls sharply.

The afterlife analysis changes the practical question. In the reference completion, cryonics produces meaningful psychological continuation in about 52.3 percent of worlds that otherwise lack the paper's broad postmortem continuation, while about 14.25 percent fall into the narrower category in which an already-continuing original subject might also be returned. Those numbers are not 'chance of rescue' versus 'chance of being dragged out of heaven.' Actual harm depends on whether a postmortem subject is active, whether the same subject returns, whether return interferes with the earlier state, whether that state is better, and whether return is chosen.

Preservation can therefore have option value even for someone uncertain about revival. A preserved patient may be able to wait for better technology, better institutional safeguards, and evidence from earlier cases. But waiting also has costs, and evidence is only as informative as the model of memory and detection behind it. A policy of 'preserve now, activate only after specified evidence' is conceptually different from automatic early revival.

No universal enrollment recommendation follows from the paper. The deeper result is methodological. A serious cryonics analysis should separate engineering success from consciousness, identity, and psychological continuation; should compare revival with the person's actual counterfactual rather than with a zero labeled 'death'; and should treat preservation, activation, consent, and learning as distinct decisions. Once those distinctions are made, readers can substitute their own metaphysical beliefs and values without losing track of what is doing the work.

About the Human Author

Rodney Malpert holds a Juris Doctor (J.D.) and Master of Arts (M.A.) from Cornell University and completed all requirements for the Ph.D. in Government except the dissertation. His Cornell graduate work also included substantial coursework in analytic philosophy. He practiced law for many years in Silicon Valley and has taught law as an adjunct professor at Southern Methodist University and Arizona State University. He achieved the rank of Senior Master with the United States Chess Federation. Beyond his formal academic and professional training, he has pursued extensive independent study across philosophy, science, history, and related fields, including hundreds of courses in The Great Courses series and sustained book and journal study. His current research uses human-AI collaboration to examine questions at the intersection of philosophy, probability, technology, and long-term human futures.

Website and current papers: RodMalpert.com

Acknowledgments, attribution, and disclosure

Rodney Malpert is the project originator and primary author. Within this collaboration he developed or pressed the psychological-continuation target alongside numerical identity; the ordinary-change and multiple-plausible-futures arguments for anchored inference; the separation of embodiment from uploading; the receiver/broadcaster concern; local versus global individuation; post-revival identity drift; and the distinction between successful and desirable revival.

His subsequent interventions made afterlife interaction central: the 90-percent good-afterlife thought experiment, substitution of readers' own beliefs, the three-stage human/postmortem/revival continuity problem, transformation during an intervening postmortem state, the conditional exit-and-return argument, memory inaccessibility on return, very long revived life as opportunity cost, early revivals as evidence, and a not-first or evidence-threshold instruction. These are attributions of contributions within the project, not assertions of historical priority. They are not empirical findings about an afterlife.

Malpert also required this revision to incorporate the 2026 global-workspace convergence result without allowing it to collapse three separate questions: whether digital computation can support access-like cognition, whether it is phenomenally conscious, and whether a conscious reconstruction would be numerically identical with the preserved person. He specifically asked that the S1/S2 computationalism distinction from the companion metaphysics work remain visible while the cryonics paper stays focused on identity rather than AI consciousness for its own sake.

ChatGPT assisted with research, argument development, calculations, and manuscript preparation. Claude supplied substantial adversarial criticism and a revised manuscript and reference program. Its contributions include symmetric survival standards across ontologies, explicit conditional partitions, the paired rescue/recall benchmark, voluntary-return analysis, a shared-firewall learning floor, the first-mover problem, joint identity/psychology bounds, and an anchoring sweep.

The final ChatGPT audit independently executed the supplied reference, retained its numerical inputs, added feasible counterfactual-dependence bounds and finite-sample learning comparisons, separated broad continuation from active-at-revival exposure, qualified consent and exit claims, and revised the interpretation of thresholds and institutional instructions. The source record and arithmetic audit are supplied separately. AI convergence is not treated as independent evidence. Malpert reports longstanding cryonics membership and a whole-body preservation arrangement; the paper does not evaluate his executed contract or forecast his personal outcome.

Appendix A. Full reference specification

Notation note. The main text deliberately uses descriptive words. This appendix returns to compact symbols for reproducibility: T = successful physical reconstruction, C = consciousness, I = same-person identity, M = meaningful psychological continuation; BL = brain-linked, HR = holistic reconstitution, ES = enduring subject, and S1 = external participant. B0 denotes broad postmortem continuation without cryonic preservation.

A.1 Joint distribution

Table A1. Ontology weights and non-S1 mechanism shares.

Ontology	Weight $w(O)$	Brain-linked (BL)	Holistic reconstitution (HR)	Enduring subject (ES)	External-participant share $s(O)$
Physicalism	0.262626	1.00	0.00	0.00	0.25
Russellian monism	0.404040	0.80	0.20	0.00	0.25
Idealism	0.232323	0.00	0.60	0.40	0.25
Dualism	0.101010	0.65	0.00	0.35	0.25

The resulting local masses are BL = 0.488636, HR = 0.165152, ES = 0.096212, and S1 = 0.250000. These are the sums of the joint cells and not independently added priors. Exact code retains the fractional ontology weights rather than the displayed rounded decimals.

A.2 Mental outcome factors

Table A2. Consciousness and continuation conditionals.

Factor	Original-brain repair	Anchored reconstruction	Digital upload
Brain-linked consciousness: physicalism	.99	.99	.55
Brain-linked consciousness: Russellian monism	.99	.99	.40
Brain-linked consciousness: dualism	.99	.99	.65
Holistic-reconstitution consciousness	.90	.90	.15
Same-person identity given consciousness (BL/HR)	.85	.80	.40
Psychological continuation given consciousness (reference)	.97	.92	.97

Table A3. Exclusive original/other attachment events.

Mechanism and event	Original-brain repair	Anchored reconstruction	Digital upload
Enduring subject: original returns (r)	.45	.45	.25
Enduring subject: other conscious subject (g)	.15	.15	.15
External participant: original returns (r)	.45	.45	.45
External participant: other conscious subject (g)	.15	.15	.15

For BL and HR, $C = c$, $I = c \cdot qI$, $M = c \cdot qM$. For ES and S1, $C = r + g$, $I = r$, $M = (r + g)qM$. The conditional Fréchet bounds for I and M are $c \max(0, qI + qM - 1)$ and $c \min(qI, qM)$ in BL/HR, and $\max(0, r + M - C)$ and $\min(r, M)$ in ES/S1. These permit identity and psychological loss to be modeled separately.

The counterfactual inputs are $b(BL) = 0.02$, $b(HR) = 0.68$, $b(ES) = 0.90$, and $b(S1) = 0.92$. The reference uses b independently of the revival outcomes within a joint cell. That independence is relaxed in Table 8 without changing the marginals.

A.3 Sensitivity settings and exact coupling bounds

Low and high attachment set every ES/S1 original/other pair to 0.10/0.10 and 0.80/0.15. The shared identity stress subtracts 0.10 from each qI. Weak P2 uses qI = 0.45 and qM = 0.55; P1-level P2 uses 0.85 and 0.97. The anchoring sweep interpolates linearly between those endpoints.

The attachment-heavy idealist case uses HR/ES shares 0.10/0.90 and sets ES original/other to 0.10/0.10 across all ontologies. It is therefore a combined structural and ES-response intervention, not a clean intervention on idealism alone. The upload-compatible case gives P3 the P1 consciousness factors, qI, and ES original return, with qM already equal. Alternative S1 shares change s uniformly. The last scenario in Table 3 changes only S1 original return for P3 from 0.45 to 0.25.

For the allocation bounds, $x(o) = P(O=o, S1)$ is constrained by $0 \leq x(o) \leq w(o)$ and $\sum x = 0.25$. The outcome is its non-S1 baseline plus $\sum x(o)[u_j - t_{oj}]$. Greedy allocation by the signed gains gives the exact linear bounds. It does not vary the accuracy of any mental response factor.

Appendix B. Counterfactual dependence and reproducibility

B.1 Feasible joint worlds

The paired-dependence audit uses ten binary atoms (B0,C,I,M). Five combinations of C,I,M are feasible because I and M imply C but not one another, and each can be paired with either value of B0. Each joint cell is constrained by total probability one and its four specified marginals. Enumerating five-atom basic feasible supports yields the vertices of this finite linear polytope.

Rescue is the linear sum of atoms satisfying not-B0 and the chosen target. ES/S1 candidate exposure is the sum satisfying B0 and I; the inclusive HR scenario adds B0 and C in HR. Weighted minimization and maximization across cell vertices yield the probability bounds. Bisection on the extrema of $S - \lambda E$ yields simultaneous bounds for S/E when the exposure lower bound is positive. This procedure uses no assumption of I/M independence and no third-party optimization package.

These bounds remain conditional on a transport of the same reference marginals into the paired worlds. They do not remove the need to model active-at-revival status, memory, interference, or institutional selection. Those are additional events, not automatically consequences of B0.

B.2 What the program reproduces

The accompanying `cryonics_model.py` and retained `claude_reference.py` use the Python standard library. Running `python cryonics_model.py --out results` regenerates the reference rows, mental-outcome conjunction bounds, all sensitivity cases, anchoring sweep, S1-coupling bounds, rescue/exposure calculations and their dependence bounds, break-even tables, shared/independent memory examples, the sequential-signal example, and finite adverse-outcome examples.

Validation checks finite probabilities, normalized allocations, exclusive attachment outcomes, and I and M no greater than C. Regression checks compare the supplied reference with independently reconstructed local totals, verify analytic probability bounds against the joint-atom enumeration, and test limiting cases and rejected invalid inputs. No numerical output is evidence that the starting philosophical or utility judgments are calibrated.

Every computed table is generated from the same machine-readable results used to typeset the manuscript. Descriptive tables and bibliographic facts are not outputs of the probability model. A separate audit note records editorial decisions and source-verification limits.

References

- Anthropic. (2026, July 6). A global workspace in language models. Interpretability research note.
- Bostrom, Nick, and Carl Shulman. (2023). Propositions Concerning Digital Minds and Society, version 1.21.
- Gurnee, Wes, et al. (2026). Verbalizable Representations Form a Global Workspace in Language Models. Transformer Circuits Thread; arXiv:2607.15495.
- McClelland, Tom. (2026). Agnosticism about artificial consciousness. *Mind & Language*, 41(4), 482-502.
- McCormack, Peter (host), with Nick Bostrom. (2026, September 24). The Peter McCormack Show #215: 'We Are Creating Superintelligence': The Existential Risk is Real.
- Agar, Nicholas. 2014. *Truly Human Enhancement: A Philosophical Defense of Limits*. MIT Press.
- Alcor Life Extension Foundation. 2026. FAQs. <https://www.alcor.org/faqs/> (accessed September 23, 2026).
- Alcor Life Extension Foundation. 2024. Cryopreservation Membership Agreement (form), sections 1.1, 4.4, and 5.2. <https://www.alcor.org/wp-content/uploads/2024/11/alcor-form-cryopreservation-agreement.pdf> (accessed September 17, 2026).
- Arrow, Kenneth J., and Anthony C. Fisher. 1974. Environmental Preservation, Uncertainty, and Irreversibility. *Quarterly Journal of Economics* 88(2): 312-319.
- Bostrom, Nick. 2003. Are You Living in a Computer Simulation? *Philosophical Quarterly* 53(211): 243-255.
- Bourget, David, and David J. Chalmers. 2023. Philosophers on Philosophy: The 2020 PhilPapers Survey. *Philosophers' Imprint* 23(11). Survey data: <https://survey2020.philpeople.org/survey/results/all>.
- Cerullo, Michael A. 2015. Uploading and Branching Identity. *Minds and Machines* 25: 17-36. DOI: 10.1007/s11023-014-9352-8.
- Chalmers, David J. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Chalmers, David J. 2010. The Singularity: A Philosophical Analysis. *Journal of Consciousness Studies* 17(9-10): 7-65.
- Chalmers, David J. 2014. Uploading: A Philosophical Analysis. In *Intelligence Unbound*, edited by Russell Blackford and Damien Broderick, 102-118. Wiley Blackwell.
- Chalmers, David J. 2016. Panpsychism and Panprotopsyism. In *Panpsychism: Contemporary Perspectives*, edited by Godehard Brüntrup and Ludwig Jaskolla, 19-47. Oxford University Press.
- Chalmers, David J. 2022. *Reality+: Virtual Worlds and the Problems of Philosophy*. W. W. Norton.
- Dainton, Barry. 2008. *The Phenomenal Self*. Oxford University Press.
- DeGrazia, David. 2005. *Human Identity and Bioethics*. Cambridge University Press.
- Diekmann, Andreas. 1985. Volunteer's Dilemma. *Journal of Conflict Resolution* 29(4): 605-610.
- German, Alexander, Enes Yağız Akdaş, Cassandra Flügel-Koch, et al. 2026. Functional Recovery of the Adult Murine Hippocampus after Cryopreservation by Vitrification. *Proceedings of the National Academy of Sciences* 123(10): e2516848123. DOI: 10.1073/pnas.2516848123.
- Gilbert, Frederic, Ian Burkhart, and Jake Morrill. 2026. The Brain-Computer Interface is Changing Me: The Psychological Trajectory of Becoming More-Than-Human with a Brain Implant. *Neuroethics* 19: 31. DOI: 10.1007/s12152-026-09656-2.
- Goff, Philip. 2017. *Consciousness and Fundamental Reality*. Oxford University Press.
- Hasker, William. 1999. *The Emergent Self*. Cornell University Press.
- Hoffman, Donald D., and Chetan Prakash. 2014. Objects of Consciousness. *Frontiers in Psychology* 5: 577. DOI: 10.3389/fpsyg.2014.00577.
- Hoffman, Donald D., Chetan Prakash, and Robert Prentner. 2023. Fusions of Consciousness. *Entropy* 25(1): 129. DOI: 10.3390/e25010129.
- Howard, Ronald A. 1966. Information Value Theory. *IEEE Transactions on Systems Science and Cybernetics* 2(1): 22-26.
- Kastrup, Bernardo. 2017. An Ontological Solution to the Mind-Body Problem. *Philosophies* 2(2): 10. DOI: 10.3390/philosophies2020010.

- Kastrup, Bernardo. 2023. AI Won't Be Conscious, and Here Is Why (A Reply to Susan Schneider). Author essay, January. <https://www.bernardokastrup.com/2023/01/ai-wont-be-conscious-and-here-is-why.html>.
- Lewis, David. 1976. Survival and Identity. In *The Identities of Persons*, edited by Amélie O. Rorty. University of California Press.
- Long, Brandon, and David Resnik. 2025. Precautionary Principle and Post-Mortem Brain Resuscitation. *Neuroethics* 18: 32. DOI: 10.1007/s12152-025-09605-5.
- Malpert, Rodney. 2026a. A Two-Level Bayesian Analysis of Reality, Consciousness, and the Physical World. Working paper, September.
- Malpert, Rodney. 2026b. A Bayesian Analysis of Life After Death. Working paper, September; version reporting a broad continuation estimate of approximately 45 percent.
- Maya-Romero, Alex, and Roberto Zoncu. 2026. Cryopreservation through Vitrification Enables Functional Recovery of the Mouse Hippocampus. *Proceedings of the National Academy of Sciences* 123(24): e2605838123. DOI: 10.1073/pnas.2605838123.
- McKenzie, Andrew T., et al. 2024. Structural Brain Preservation: A Potential Bridge to Future Medical Technologies. *Frontiers in Medical Technology* 6: 1400615. DOI: 10.3389/fmedt.2024.1400615.
- Merkle, Ralph C. 1992. The Technical Feasibility of Cryonics. *Medical Hypotheses* 39(1): 6-16. DOI: 10.1016/0306-9877(92)90133-W.
- Miller, Allen C. 1975. The Value of Sequential Information. *Management Science* 22(1): 1-11.
- Minerva, Francesca. 2018. *The Ethics of Cryonics: Is It Immoral to Be Immortal?* Palgrave Macmillan.
- Moen, Ole Martin. 2015. The Case for Cryonics. *Journal of Medical Ethics* 41(8): 677-681. DOI: 10.1136/medethics-2015-102715.
- Nozick, Robert. 1981. *Philosophical Explanations*. Harvard University Press.
- Olson, Eric T. 1997. *The Human Animal: Personal Identity Without Psychology*. Oxford University Press.
- Parfit, Derek. 1984. *Reasons and Persons*. Clarendon Press.
- Parfit, Derek. 2012. We Are Not Human Beings. *Philosophy* 87(1): 5-28.
- Paul, L. A. 2014. *Transformative Experience*. Oxford University Press.
- Prentner, Robert, and Donald D. Hoffman. 2024. Interfacing Consciousness. *Frontiers in Psychology* 15: 1429376. DOI: 10.3389/fpsyg.2024.1429376.
- Schechtman, Marya. 1996. *The Constitution of Selves*. Cornell University Press.
- Seth, Anil K. 2025. Conscious Artificial Intelligence and Biological Naturalism. *Behavioral and Brain Sciences*, target article, accepted manuscript. DOI: 10.1017/S0140525X25000032.
- Shani, Itay. 2015. Cosmopsychism: A Holistic Approach to the Metaphysics of Experience. *Philosophical Papers* 44(3): 389-437.
- Shaw, David. 2009. Cryoethics: Seeking Life after Death. *Bioethics* 23(9): 515-521.
- Shoemaker, Sydney. 1984. Personal Identity: A Materialist's Account. In *Personal Identity*, by Sydney Shoemaker and Richard Swinburne, 67-132. Blackwell.
- Sotala, Kaj, and Harri Valpola. 2012. Coalescing Minds: Brain Uploading-Related Group Mind Scenarios. *International Journal of Machine Consciousness* 4(1): 293-312.
- Stecker, Mark M., et al. 2001a. Deep Hypothermic Circulatory Arrest: I. Effects of Cooling on Electroencephalogram and Evoked Potentials. *Annals of Thoracic Surgery* 71(1): 14-21.
- Stecker, Mark M., et al. 2001b. Deep Hypothermic Circulatory Arrest: II. Changes in Electroencephalogram and Evoked Potentials During Rewarming. *Annals of Thoracic Surgery* 71(1): 22-28.
- Story, Daniel. 2024. Cryonics: Traps and Transformations. *Bioethics* 38(4): 351-355. DOI: 10.1111/bioe.13277.
- Swinburne, Richard. 1997. *The Evolution of the Soul*, revised edition. Clarendon Press.
- Tononi, Giulio, and Christof Koch. 2015. Consciousness: Here, There and Everywhere? *Philosophical Transactions of the Royal Society B* 370: 20140167.

- Unger, Peter. 1990. *Identity, Consciousness and Value*. Oxford University Press.
- van Inwagen, Peter. 1990. *Material Beings*. Cornell University Press.
- Zimmerman, Dean. 2010. From Property Dualism to Substance Dualism. *Proceedings of the Aristotelian Society*, Supplementary Volume 84: 119-150.

Appendix C. Preservation duties, AI capability, and the conditions of revival

The Bostrom-Yampolskiy interview connects uploading, preservation, and the welfare of digital minds. Applied to cryonics, its strongest contribution is to separate the power to restore a mind from the responsibility to restore it faithfully and treat it well. This is an extension of the interview's discussion to the present decision problem, not a claim that the participants endorse this paper's numerical model.

C.1 Preservation has two different objects

Keeping an original brain can protect a possible continuing subject under some theories of identity. It can also preserve evidence needed to constrain reconstruction. Those benefits differ. Even if a conscious upload is a meaningful psychological continuer, it may be valuable to retain the biological original while questions about strict identity or fidelity remain unresolved.

Conversely, a backup of a digital state is not a demonstrated analogue of every cryonics procedure. Biological preservation may leave damage and ambiguity requiring repair or inference. A statement that digital information can be retained cheaply does not establish that a particular preserved brain contains everything required to restore its person. The reference model's strong surviving anchor must remain explicit.

C.2 What successful early revival could establish

Evidential target	Useful observation	Important limit
Function and communication	Stable capacities, comprehension, adaptive behavior	Function alone is not proof of phenomenality
Psychological fidelity	Independent private records, withheld autobiographical tests, error accounting	Answers installed during reconstruction are not independent confirmation
Present welfare and agency	Persistent preferences, communication, opportunities to refuse and revise choices	Apparent satisfaction can be engineered or constrained
Strict numerical identity	Theory-sensitive continuity evidence and discriminating predictions where available	Inherited memories can produce sincere identity reports in distinct successors
An intervening afterlife	Independently checkable information with ordinary routes controlled	Sincerity and vividness do not establish an external source

These limits do not make empirical learning futile. They identify its target. A series of coherent revivals could improve confidence in fidelity and welfare more than in an identity theory whose alternatives predict the same reports. Selection of better-preserved patients

and changes in procedures would also limit transfer to other cohorts.

C.3 Technical success is not the full probability of benefit

The Section 8 scores condition on successful construction and retain their stated worldview weights. They do not estimate the probability that a present patient is preserved adequately, remains in custody, becomes technically repairable, is selected for revival, and receives a life they would endorse. AI progress could improve the repair step while worsening the selection, autonomy, or welfare steps.

A complete patient-outcome calculation would condition on a world state containing technology and governance, then model preservation, institutional continuity, selection, construction, consciousness, fidelity, and welfare. The chain rule permits successive conditional factors; it does not permit multiplying unrelated unconditional estimates. A highly capable benevolent system, a capable indifferent system, and a capable coercive system can have similar engineering ability and radically different patient outcomes.

This also explains why the AI-risk headline cannot be subtracted from cryonics survival scores. Civilizational extinction, severe dystopia, and an individual's desired continuation are different events with overlapping but unmodeled dependencies. A future that preserves humanity collectively might still disregard one patient's wishes. A repaired original might survive in circumstances the patient would reject.

C.4 Custody can create continuing moral responsibilities

An institution accepting custody takes responsibility for interests that may become actionable long after the original agreement. Normative duties can include maintaining material and records, preserving provenance, representing the patient's instructions, distinguishing known facts from reconstruction, and protecting any conscious subject who results. These are ethical proposals; their legal enforceability is a separate question.

Credible commitment requires more than a promise of benevolence. Who represents the patient after organizational succession? Who may authorize activation, copying, enhancement, or merger? Who can challenge a decision? What happens when funding, ownership, or control changes? A trustworthy arrangement needs authority, incentives, and continuity of records appropriate to those questions.

Preservation can also transfer control. The option value of "preserve now, decide later" belongs meaningfully to the patient only insofar as later decision-makers can wait, learn, and honor relevant preferences. That consideration does not remove the costs of waiting or prove that every delay is beneficial. It places governance inside the option-value argument.

C.5 Restoration, compensation, and transformation

Restoration does not automatically discharge responsibility for avoidable damage, coercive experimentation, or unwanted alteration. Compensation raises a further identity question: benefiting a psychological continuer need not conclusively compensate the original experiencer. Still, the resulting conscious person's interests matter even if strict identity remains disputed.

Enhancement should be evaluated separately from initial restoration. Bostrom's concern about whether favorable dispositions survive radical change has a direct analogue in patient care: a technically successful revival can be followed by changes that weaken the person's attachment to earlier relationships or projects. Some changes may be freely endorsed. The relevant protection is meaningful comprehension and choice, not an assumption that all psychological development is harmful.

C.6 What cryonics adds if an operator already preserves us

Let B denote qualifying continuation without cryonic preservation and R the desired continuation through revival. The marginal benefit is not $P(R)$ alone. If an operator already retains sufficient states and restores subjects independently, cryonics may add little technical preservation value. It might still communicate preferences, preserve in-world evidence, or affect an operator's policy. Those are further hypotheses requiring their own support.

If operator restoration and cryonic revival coexist, there may be two continuers rather than one displaced subject. If they compete for the same original subject, interference becomes relevant. Neither arrangement follows simply from the word "simulation." A knowledgeable benevolent operator could reduce harmful return, but benevolence without knowledge does not establish that protection.

C.7 Numerical relation to the companion papers

The 43.87% B0 reference in this paper is retained as its own counterfactual mixture. The current life-after-death benchmark is 46.56%; it uses different local weights and an explicit empirical update. Its generated-mind extension gives 51.20% if that branch's continuation conditional is 50%. These values are not interchangeable forecasts of the same standardized patient population.

With the cryonics psychological-continuation marginal $m = 81.5\%$, the general bound $\max(0, m-b)$ to $\min(m, 1-b)$ gives a useful comparison. At $b = 43.87\%$, rescue can lie between approximately 37.63% and 56.13%; at $b = 46.56\%$, between 34.94% and 53.44%. These are broad marginal-only bounds, not replacements for the much narrower mechanism-conditioned bounds in Table 8. Raising the no-preservation continuation marginal by itself does not identify the new coupling or a revised harm estimate.

Accordingly, the existing pathway scores and mechanism-conditioned rescue/exposure tables remain reference scenarios. The interview supplies no independently calibrated replacement inputs. The improvement is a more complete decision structure: preservation, restoration, original-subject survival, and flourishing require different achievements.

Sources

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), Chalmers (2010; 2014), and the cryonics and decision-theory sources cited in Sections 2-3 and 9-12 provide the underlying continuity and option-value arguments. Bostrom and Shulman (2023), Propositions Concerning Digital Minds and Society, version 1.21, discusses welfare and consent for digital subjects. <https://nickbostrom.com/propositions.pdf>. Patient-representation proposals and the cross-paper calculations here are our extensions.