

Cryonics and personal survival

RodMalpert.com

Executive summary

Cryonics raises distinct questions about physical reconstruction, consciousness, strict identity, meaningful psychological continuation, and whether revival improves the person's outcome. The paper compares successful future constructions; it does not estimate a present patient's probability of reaching that stage.

Pathway	Same-person identity	Meaningful psychological continuation
Original-brain repair	69.3%	81.5%
Strongly anchored reconstruction	66.2%	77.3%
Conventional digital upload	24.1%	43.7%

These reference scores are philosophical scenario outputs, not observed rates. They remain unchanged. The relative advantage of repair depends on retained person-specific information and on assumptions about which implementations can support consciousness and identity.

The Bostrom-Yampolskiy interview adds a sharper account of preservation and custody. Retaining an original can protect a possible subject and preserve evidence for better reconstruction. Neither a backup nor a successor's sincere identity report automatically proves original-subject survival.

AI could improve technical restoration while worsening selection, autonomy, or welfare. A complete patient-outcome model must connect those events conditionally. Institutional duties, patient representation, reliable records, and separate consent for restoration, enhancement, acceleration, and merger are therefore central.

The no-preservation counterfactual is not automatically zero. The paper retains its own 43.87% broad-continuation reference mixture. The companion survival benchmark is 46.56%, with generated-preservation sensitivity extending above one half. The different mixtures cannot be reconciled by subtracting or copying headline percentages.

Preservation may retain valuable choices, but only if future custodians can learn and respect relevant preferences. Existing operator backups could reduce cryonics' technical added value; harmful interference and parallel continuation are different possibilities. The interview provides no calibrated probability for either. No universal enrollment or activation recommendation follows.

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.