

Summary of the Framework and Proposals

Precautionary Framework

A *sentient being* (in the sense relevant to the present framework) is a system with the capacity to have valenced experiences, such as experiences of pain and pleasure.

A system *S* is a *sentience candidate* if there is an evidence base that (a) implies a realistic possibility of sentience in *S* that it would be irresponsible to ignore when making policy decisions that will affect *S*, and (b) is rich enough to allow the identification of welfare risks and the design and assessment of precautions.

A system *S* is an *investigation priority* if it falls short of the requirements for sentience candidature, yet (a) further investigation could plausibly lead to the recognition of *S* as a sentience candidate and (b) *S* is affected by human activity in ways that may call for precautions if *S* were a sentience candidate.

Framework Principle 1. *A duty to avoid causing gratuitous suffering.* We ought, at minimum, to avoid causing gratuitous suffering to sentient beings either intentionally or through recklessness/negligence. Suffering is not gratuitous if it occurs in the course of a defensible activity despite proportionate attempts to prevent it. Suffering is gratuitous if the activity is indefensible or the precautions taken fall short of what is proportionate.

Framework Principle 2. *Sentience candidature can warrant precautions.* If *S* is a sentience candidate, then it is reckless/negligent to make decisions that create risks of suffering for *S* without considering the question of what precautions are proportionate to those risks. Reasonable disagreement about proportionality is to be expected, but we ought to reach a policy decision rather than leaving the matter unresolved indefinitely.

Framework Principle 3. *Assessments of proportionality should be informed, democratic, and inclusive.* To reach decisions, we should use informed and

inclusive democratic processes. These decisions should be revisited periodically and whenever significant new evidence emerges.

An example of an informed, inclusive, democratic process is a citizens' panel or assembly that assesses the proportionality of possible responses by debating their permissibility-in-principle, adequacy, reasonable necessity, and consistency (the 'PARC' tests).

Proposals about Specific Cases

The following should be read as proposals for discussion and debate. There is a case for regarding the proposed measures as proportionate, so they merit consideration by an informed, inclusive democratic process, but my aim is not to pre-judge the outcome of such a process.

People with Disorders of Consciousness

Proposal 1. *Assume sentient (specific).* Any signs that would be interpreted as signs of pain or distress in a conscious patient should still be so interpreted in a patient with any disorder of consciousness, and pain relief administered. This treatment should not be reserved for a subset of patients, e.g. those diagnosed as minimally conscious.

Proposal 2. *Assume sentient (general).* A patient with a prolonged disorder of consciousness should not be assumed incapable of experience when an important clinical decision is made. All clinical decisions should consider the patient's best interests as comprehensively as possible, working on the precautionary assumption that there is a realistic possibility of valenced experience and a continuing interest in avoiding suffering and in achieving well-being, but without taking this assumption to have implications regarding prognosis.

Proposal 3. *Avoid line-drawing (general).* The diagnostic categories of PVS, MCS–, and MCS+ should be phased out in therapeutic and legal contexts. Clinicians should work with the broader category of 'prolonged disorder of consciousness' (PDOC) and with profiles of individual patients, tailoring care to the patient's individual needs.

Proposal 4. *Avoid line-drawing (specific).* Decisions about withdrawing life-sustaining treatment should be based on comprehensive assessment of the patient's best interests, sensitive to the details of their case, and not on the PVS/MCS distinction.

Proposal 5. *The need for more humane options.* Methods of hastening death other than through withdrawal of clinically assisted nutrition and hydration (CANH) are needed. At minimum, clinicians should not face any risk of punishment for administering large doses of sedatives or analgesics after a decision to withdraw CANH has been made.

Proposal 6. *Waiting for more information.* If a humane method of hastening death becomes available, there will be no strong rationale for withdrawing life-sustaining treatment within days after injury, before the patient's condition has stabilized. Given the great uncertainty about prospects for recovery at this stage, the default approach should be to delay the decision until the patient's condition has stabilized.

Human Fetuses and Embryos

Proposal 7. *Sentience and abortion.* The point at which a human fetus becomes sentient is not the point at which abortion becomes morally impermissible. We should separate these issues. The ethics of abortion depend primarily on questions of personhood and bodily autonomy, not on questions of sentience.

Proposal 8. *Human sentience candidature begins early.* Human fetuses are sentience candidates from the beginning of the second trimester. This line may move as new evidence emerges, but it should always track the earliest scientifically credible, evidence-based estimate.

Proposal 9. *Fetal pain relief (in therapeutic contexts).* Direct fetal anaesthesia and pain relief should be considered whenever therapeutic fetal surgery is performed. The public should be involved in discussions about general norms of medical practice (not specific decisions), and data regarding current practices should be collected and published to allow such discussions.

Proposal 10. *Fetal pain relief (in the context of abortion).* Clinicians need to communicate uncertainty about fetal sentience honestly to patients. In some

cases, fetal pain relief may be appropriate. Deliberative processes for setting clinical norms must give appropriate weight to the voices of women.

Proposal 11. *Sentience and the 14-day rule.* If the main goal of setting a legal time limit on human embryo research were that of prohibiting research on human sentience candidates, a significant liberalization of the current 14-day rule would still be proportionate to that goal. However, the issue raises deep value conflicts that have little to do with sentience.

Human Neural Organoids

Proposal 12. *Brainstem rule.* If a neural organoid develops or innervates a functioning brainstem (including the midbrain) that regulates arousal and leads to sleep-wake cycles, then it is a sentience candidate. An artificial functional equivalent of a brainstem would also suffice.

Proposal 13. *Targeted bans.* If organoid research leads to the creation of organoids that are sentience candidates, a moratorium (time-limited ban) or indefinite ban on the creation of this particular type of organoid may be an appropriate response. Bans should avoid indiscriminate targeting of all organoid research.

Proposal 14. *Ethical review.* When a neural organoid is a sentience candidate, research on it, if permitted at all, should be subject to ethical review and harm-benefit analysis, modelled on existing frameworks for regulating research on sentient animals.

Other Animals

Proposal 15. *All adult vertebrates are sentience candidates.* Debates about proportionality are warranted in cases where human activities create risks of suffering to any adult vertebrate animal. Further investigation concerning sentience candidature in vertebrates should focus on juvenile/larval stages.

Proposal 16. *Sentience is neither intelligence nor brain size.* We should be aware of the possibility of decouplings between intelligence, brain size, and sentience in the animal kingdom. Precautions to safeguard animal welfare

should be driven by markers of sentience, not by markers of intelligence or by brain size.

Proposal 17. *Sentience is not pain.* Although there are pragmatic reasons for the focus on pain in debates about animal sentience, we must be open to the possibility that the class of sentient animals is bigger than the class of animals that feel pain. Other forms of evidence can make an animal a sentience candidate, such as evidence of sophisticated forms of learning, attention, working memory, and planning.

Proposal 18. *Some invertebrates are sentience candidates.* Coleoid cephalopod molluscs, decapod crustaceans of the suborder Pleocyemata, and insects (all when in the adult stage) are sentience candidates. Debates about proportionality are warranted in cases where human activities create risks of suffering to these animals. Decapod crustaceans of the suborder Dendrobranchiata, insect larvae, spiders, gastropods, and nematode worms are investigation priorities.

Proposal 19. *Codes of good practice and licensing.* There should be a licensing scheme for companies attempting to farm sentience candidates or investigation priorities for which no welfare regulations yet exist (such as insects). Obtaining a license should be dependent on signing up to (and, where necessary, funding research leading to) a code of good practice concerning animal welfare.

Proposal 20. *Octopus farming.* It is very unlikely that octopus farming can meet reasonable expectations regarding welfare and humane slaughter. It would be proportionate to ban octopus farming.

Proposal 21. *Towards humane slaughter.* When an animal is a sentience candidate, it is proportionate to ban slaughter methods that needlessly risk extending and intensifying the suffering associated with dying, such as boiling animals alive without prior stunning.

Artificial Intelligence (AI)

Proposal 22. *Sentience is not intelligence (II).* We should be aware of the possibility of a substantial decoupling between intelligence and sentience in the AI domain. Precautions to manage risks of suffering should be driven by markers of sentience, not markers of intelligence. For example, emulations of

animal brains could achieve sentience without necessarily displaying impressive intelligence.

Proposal 23. *The gaming problem.* For any set of criteria for sentience candidature, we need to be aware of the risk of the AI system or its designer learning (implicitly or explicitly) that they are regarded as criteria, leading to gaming of the criteria. We need to discount markers we have reason to think may have been gamed.

Proposal 24. *Deep computational markers.* We can use computational functionalist theories (such as the global workspace theory and the perceptual reality monitoring theory) as sources of markers of sentience. If we find signs that an AI system, even if not deliberately equipped with such features, has implicitly learned ways of recreating them, this should lead us to regard it as a sentience candidate.

Proposal 25. *The run-ahead principle.* At any given time, measures to regulate the development of sentient AI should run ahead of what would be proportionate to the risks posed by current technology, considering also the risks posed by credible future trajectories.

Proposal 26. *Codes of good practice and licensing (II).* There should be a licensing scheme for companies attempting to create artificial sentience candidates, or whose work creates even a small risk of doing so, even if this is not an explicit aim. Obtaining a license should be dependent on signing up to (and, where necessary, funding the creation of) a code of good practice for this type of work that includes norms of transparency.