

The Modalities and Practicalities of Hell

by The Legendary Dragon Alephwyr

Philosophy offers many ontological and metaphysical speculations and doctrines. In our current, materialist age, it can be hard not to see these as merely metaphors or impartial glimpses of more mundane truths. This is unfair to the concepts in question, since they once stood quite handily on their own, as part of larger networks of knowledge and thought that have since been lost to all but a small number of academics, religious experts, and diligent hobbyists at best. However, unfair or not, it is in the nature of philosophical ideas to be subsumed into new networks of knowledge and new paradigms of thought, and it would be remiss of us not to explore the potential implications of this as it pertains to some of the more significant metaphysical doctrines of the world. Old ideas are never abandoned, they are appropriated when useful. To understand what ideas may be useful in the future, it is simple enough to look at what ideas have been useful in the past.

Some of the most widespread and socially influential doctrines in the world are doctrines of hell. Most major religions of the world have some doctrine of hell, including Christianity, Buddhism, Islam, and Hinduism. While these doctrines differ in their descriptions of the intensity, nature, causes, duration, and permanence of hell, they typically hold that hell is a literal place where people are punished for transgressions or behaviors. In some cases hell is purely punitive, and in others it is rehabilitative, but in most cases it is real. The doctrine of hell's reality, its cultural potency, and its potential implementability in the context of future technologies, all make hell doctrines a concept to watch as the boundaries between metaphysical and mundane physical interpretations blur due to improvements in science and technology and resulting cultural shifts.

By necessity, discussions about hell in any context are likely to be very morally loaded. I am not here to offer my moral opinion on any particular issue, but to explore a variety of potential technologies that might correspond to a literal hell in the future, as well as what the utility or liability of these technologies might be in the future. Lastly, to the extent of my ability I will try to discuss the potential overlap of these techniques, technologies, and their utilities and liabilities, with more traditional hell concepts.

Chapter 1

Synthetic consciousness

Needless to say, the limitations of biology preclude the possibility of creating a hell in this medium in all but the most metaphorical sense. Therefore, any implementation of a hell requires the discovery and implementation of synthetic consciousness, whether this is computerized consciousness, or some other sort of machine consciousness that likely interfaces with a computer. The particulars of such a technology are of course unknowable as of this time, but if consciousness can interface with a computer, that makes certain things relatively knowable or predictable. For the sake of argument, I will also be assuming (rightly or wrongly) that future synthetic consciousness is capable of modularity and hot-swappability, whether because it is computerized or for other unknown reasons.

Computer processes can be forked, merged, throttled up or down, and bifurcated into components without any process cessation. If all of these things are also possible with consciousness, then a number of important things become possible:

1. The duplication of consciousness
2. The subsumption of consciousness into other consciousness
3. The subjective increase or decrease in rate of time flow
4. The removal of unwanted aspects of consciousness
5. The division of different aspects of consciousness into different discrete consciousnesses

In particular, 3 is very important to the creation of an artificial hell. 2 and 4 have more relevance to doctrines of reincarnation and corrective hell, respectively. 5 has potential relevance to some concepts of the soul as expressed in Gnostic, Hermetic, and Pythagorean philosophy.

All doctrines of hell concern very large time spans. Some allow for or stipulate eventual liberation from hell via reincarnation or ascension. In both of these cases, hell is seen to have a purifying effect on the soul. If consciousness increasingly comes to be seen as the soul, then the

removal of unwanted components of consciousness (and potential replacement by other components) will be necessary techniques or technologies for implementing corrective or purifying versions of hell. Other doctrines suggest that the soul is separated from the body and/or the spirit, the latter of which do not necessarily desist, either literally or in a broader metaphysical context. I will refer to this general practice as de-husking.

There are, broadly, four kinds of hell I can think of. Full consciousness insertion, in which an unaltered consciousness is inserted in its entirety into a hell. Dehusk and fire, in which a wanted portion of a person's consciousness is dehusked and the remainder is cast into hell. Dehusk in fire, in which a full and unaltered consciousness is inserted into hell, then the wanted portion is dehusked and the remainder deleted. And Dehusk in fire and fire, in which case the wanted portion is dehusked in hell, and the unwanted portion remains in hell. The term "fire" is used as a stand-in or abstraction for whatever particulars make the hell, hell, but may also consequently refer to the actual subjective perception of fire.

Additional technological abilities that would be very helpful if possible are the ability to encrypt consciousness or its technological interface, and digital rights management style copying prevention mechanisms. The ability to delete or modify data is very important, and for reasons that will be discussed later this importance makes absolute and sole sovereign control over deletion and modification important and necessary preconditions to any artificial hell, as well as important preconditions to various other ends.

Chapter 2

Utilities and liabilities

It is generally sound practice to put warnings as close to the beginning of a document as one can. Therefore, this second chapter concerns the utilities and liabilities of an artificial hell as I am able to conceive of them, under different configurations and circumstances.

To begin with, the creation of an artificial hell in general constitutes a red line roughly analogous to the use of nuclear weapons. Nuclear weapons represent an existential threat, therefore the logic of nuclear weapons has developed into a generalized deterrence for the most part, barring a few exceptions like so-called “tactical” nukes. Even the testing of nuclear weapons can exacerbate tensions.

The creation of an artificial hell, however, represents a worse-than-existential threat. Existential threats constitute the cessation of all future utility, but an artificial hell represents an instantiation of extremely large negative utility. Therefore we should rationally expect both a much more significant deterrence effect from an artificial hell held *in potentia*, and a much stronger disincentive against actually deploying or using artificial hell technology. Human beings have done ok with nukes, even in cases where they were in the hands of ostensible lunatics, in large part due to the universally known logic of deterrence and mutually assured destruction. As long as the logic of the artificial hell is similarly widely understood, we should probably be able to expect the same kind of outcome.

However, some standard cases and edge cases are worth exploring in more depth in order to understand the particulars of this unique situation. To begin with, minus an especially strong humanitarian impulse, the threat of hell is only credible deterrence to the extent the verisimilitude of virtual environments is absolute, or realistically expected to become absolute. This is because, if one can rate the likelihood one is in a virtual environment at 0, there is no risk of personally experiencing a hell. The effective deterrence of the artificial hell thus scales upward over time as technology

improves. Therefore, the logic of the artificial hell creates a deterrence effect as the subjective perception of historical and technological progress approaches the practical potential for an artificial hell, even before these reach the time at which the first artificial hell is “created”. Like nukes, one expects the primary use case of the artificial hell will be deterring the deployment of any other artificial hell.

Another important consideration is what level of consciousness warrants protection. Would the mass torture of animals, for instance, elicit the same response? What about levels of artificial consciousness that lack even the cognition of animals? Ultimately, since these questions are liable to be somewhat subjective or at least to find varying entrenched opinions one way or the other about them, the deterrence effect will probably cover most forms of consciousness. The abortion debate is instructive: while there is a measurable point prior to which a fetus has substantive consciousness, people still blow up abortion clinics in response to the destruction of human life and potential.

If the artificial hell were to have a use case besides deterrence against the deployment of any other artificial hell, it would, however, have to limit either the scope of torture in terms of time or the scope of consciousness subjected to torture. It would also have to do so in an accountable manner. Tactical nukes are theoretically usable without triggering mutually assured destruction because they are separate from larger yield nukes in an accountable and observable manner, from the mechanisms and logistics of their deployment to the scope and potential use cases of their destruction.

Limiting time is less ambiguous and less debatable than limiting consciousness, but limiting consciousness would be an interesting approach both for its relation to various metaphysical doctrines and for the sort of divisions it would produce. For instance, assuming all religious people became religious in a weird, future way where an artificial hell satisfied for them the conditions of traditional hell concepts, the belief in essential human nature would make even very limited, human-derived consciousnesses valid targets of torture. This could allow for a compromise between secular and religious moralists: since the smallest human-derived unit of consciousness would still be human

to a religious person, this would function as an inversion of anti-abortion logic and allow a case in which a secular person had limited objections to indefinite or long-term torture, while satisfying the moral intuition of the religious that such torture is indicated.

However, in practice there are many religions and many different moral intuitions even within the context of a single religious framework. Therefore, who some might consider worthy of long-term or indefinite torture, others might consider unworthy, creating a significant mess. It would thus be much safer, even in cases of widespread agreement, not to torture anyone indefinitely, and to restrict torture in an accountable way. Even so, there are different gradations of risk, or levels of risk, and it is worth being honest about them. Attached is a picture that best summarizes my non-expert expectation of the different risk levels associated with creating different types of artificial hell.

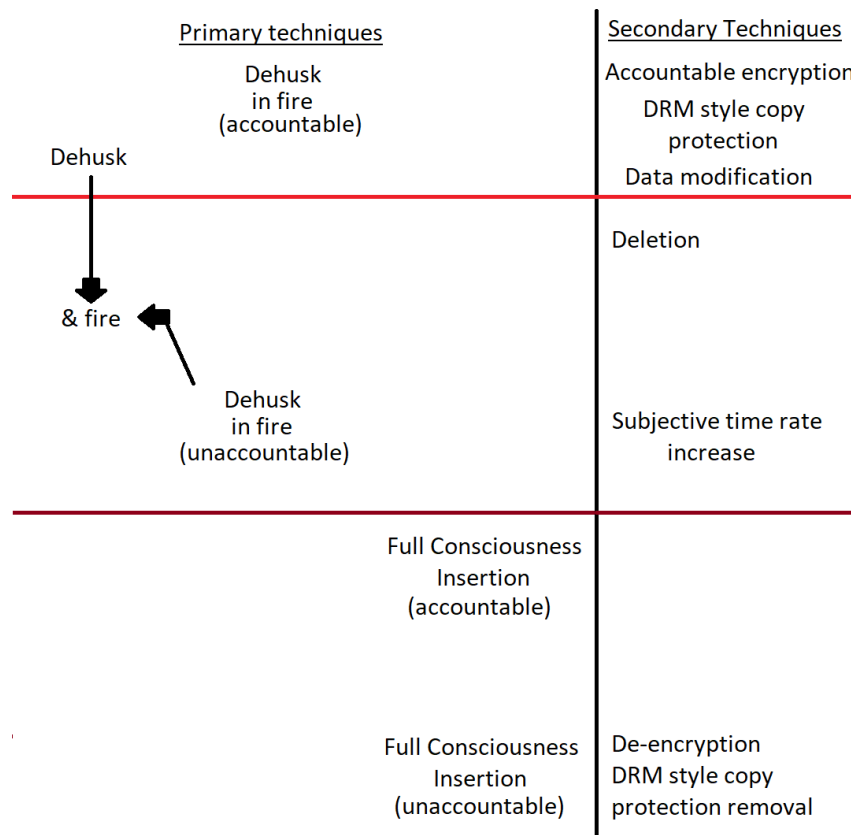


Figure 1

This picture represents, at best, a very rough guess as to the level of risk associated with each technology and technique, in increasing order of risk as you go down. It is included mostly for the sake of having a picture.

De-encryption and DRM style copy protection removal are listed as extremely high risk for several reasons. First, they allow the duplication or modification of consciousnesses that have been deemed evil enough to deserve hell. Second, they undermine the assumption of absolute sovereignty necessary to justify hell: A supreme authority is needed for such rationalization, therefore allowing others besides the supreme authority access to the damned is effectively saying "I extend my authority to anyone in this matter". This undermines the rationale for hell and hence the credibility of the authority responsible for hell.

Chapter 3

Rationalizations and justifications

To a certain extent, we can expect these to just be the same rationalizations and justifications used for all torture, which don't need to be elucidated especially due to their ubiquity. However, some religious or pseudo-religious justifications will also certainly be provided. Some sects of Christianity hold that the resurrection will occur on earth. The belief that hell could also occur on earth in the case of the damned is an intuitive leap either directly from here, or from the same sort of evolution in logic that birthed this concept.

The concept of eternity as the permanence of anything established in time rather than literal eternal duration also offers a back door to the creation of artificial hell: the desire to etch the suffering of some hated people permanently into time. Also possible is the use of artificial hell in a ritualistic manner to express intentionality, with the "sending to hell" as an act expressing alignment with God's or karmic will: Perhaps God or Karma are thought to "take over" after the initial act.

One other possibility is that as lifespans grow increasingly large, the palliative of death will cease to assuage the injured. It will no longer be possible to say "all are equal in death" when some people seemingly cannot die. This is likely to be interpreted as the avoidance of God's will and God's law, even temporarily. Therefore, the creation of artificial hell may be justified on the basis that it brings material reality more closely into alignment with God's will by minimizing its circumvention.

Regardless of what rationalizations prevail, wherever there is an injury to be inflicted, there is a risk to be managed. Risks tend to be managed by the same mechanisms through which they are summoned. Therefore, even if there is absolutely no intrinsic will at all to create or maintain an artificial hell, the artificial hell is certain to become an important part of the deterrence of future societies simply because it is possible. It is my hope that this document will somehow help start a conversation about these concepts that establishes an effective logic and equilibrium of mutual deterrence. Barring absolute epistemological certainty of sole superpower, this is necessary.