Chapter 8A
Robin Hanson on Miller's "Some Economic Incentives Facing a Business that Might Bring About a Technological Singularity"

James Miller imagines a public firm whose product is an artificial intelligence (AI). While this AI device is assumed to become the central component of a vast new economy, this firm does not sell one small part of such a system, nor does it attempt to make a small improvement to a prior version. Miller instead imagines a single public firm developing the entire system in one go. Furthermore, if this firm succeeds, it succeeds so quickly that there is no chance for others to react – the world is remade overnight.

Miller then focuses on a set of extreme scenarios where AIs "destroy the value of money". He gives examples: "mankind has been exterminated, … the new powers that be redistribute wealth independent of pre-existing property rights, …, [or] all sentient beings are merged into a single consciousness". Miller's main point in the paper is that a firm's share prices estimate its financial returns conditional on money still having value, yet we care overall about unconditional estimates. This can lead such an AI firm to make socially undesirable investment choices.

This is all true, but is only as useful as the assumptions on which it is based. Miller's chosen assumptions seem to me quite extreme, and quite unlikely. I would have been much more interested to see Miller identify market failures under less extreme circumstances.

By the way, an ambitious high-risk AI project seems more likely to be undertaken by a private firm, vs. a public firm. In the US, private firms accounted for 54.5 % of aggregate non-residential fixed investment in 2007, and they seem 3.5 times more responsive to changes in investment opportunities.[6] Public firms mostly only undertake the sorts of investments that can give poorly informed stock speculators reasonable confidence of good returns. Public firms leave subtler opportunities to private firms. Since 83.2 % of private firms are managed by a controlling shareholder, a private firm would likely, when choosing AI strategies, consider scenarios where the value of money is destroyed. So to the extent that public firm neglect of such scenarios is a problem, we might prefer private firms to do ambitious AI research.

[6] http://papers.nber.org/papers/w17394