

The Social Science of Computerized Brains
Book Review of *The Age of Em: Work, Love, and Life When Robots Rule the Earth*
by Robin Hanson (Oxford University Press, 2016)

Seth D. Baum
Global Catastrophic Risk Institute
<http://sethbaum.com> * <http://gcri.org>

Futures 90:61-63 (2017), [doi:10.1016/j.futures.2017.03.005](https://doi.org/10.1016/j.futures.2017.03.005)
This version 12 June 2017

Abstract

A new book by Robin Hanson, *The Age of Em: Work, Love, and Life When Robots Rule the Earth*, is reviewed. *The Age of Em* describes a future scenario in which human minds are uploaded into computers, becoming emulations or “ems”. In the scenario, ems take over the global economy by running on fast computers and copying themselves to multitask. The book’s core methodology is the application of current social science to this future scenario, a welcome change from the usual perspectives from physical science, computer science, and philosophy. However, in giving a wide tour of the em world, the book sometimes gets bogged down in details. Furthermore, while the book claims that the em takeover scenario would be a good thing for the world and thus should be pursued, its argument is unpersuasive. That said, the book offers by far the most detailed description of the em world available, and its scenario offers a rich baseline for future study of this important topic.

Keywords: artificial intelligence; brain emulation; scenario analysis; economics

1. Introduction

The Age of Em: Work, Love, and Life When Robots Rule the Earth by Robin Hanson discusses the possibility of “uploading” human minds into computers. An uploaded mind or “em” (short for “emulation”) is able to think more or less like a regular biological human, but with properties of computer programs and data. Ems can be copied *ad infinitum* and sped up or slowed down by being placed on faster or slower hardware. These capabilities create radical possibilities, which is the book’s focus. The book is by far the most detailed description of the em world available. Hanson, an economist with a background in physics and artificial intelligence (AI), brings a distinctive and valuable perspective to this important topic.

The Age of Em describes one em world scenario in detail, giving only brief attention towards the end of the book (Ch. 28) to variations of this scenario. The chosen scenario is rated as the most likely one, given that uploading occurs (p.34-37). The scenario is intended as a baseline that is representative of many possible scenarios; interested analysts can introduce variations to produce still other scenarios.

The book states that “Such straightforward scenario construction seems taboo among many professional futurists today”, who instead favor creating “a modest number of scenarios to cover a wide range of possibilities across key axes of uncertainty and disagreement, with each scenario being story-like, archetypical in representing clusters of relevant driving forces, and describing equilibrium rather than transitory situations” (p.35), citing Schoemaker (1995). It is certainly true that some futures studies feature sets of story-like equilibrium scenarios, and it may well be true

that some futurists disprefer presenting a single scenario, covering transition periods. However, in my experience, futurists are open to a wide range of approaches, and I see no reason why they should reject that of *The Age of Em*.

The bulk of the book consists of descriptions of the world in which the ems live. The scope is encyclopedic, with attention to everything from the social relations among ems to the resistance of em buildings to earthquakes. Indeed, the book may be best treated as an encyclopedia of a particular em scenario. Reading it straight through, it can get bogged down in details. Topics are also often presented in an unintuitive order. For example, the initial creation of the em world appears in subsections of chapters 12 and 28, instead of in its logical place at the beginning of the book. A more story-like structure would make it more readable and may reduce certain cognitive biases (Steinhardt and Shapiro 2015). However, if you are wondering about what some aspect of an em world might look like, odds are good that a description can be found somewhere in the book.

A central premise of the book's scenario is that the em world is organized around market activity. Ems exist because they provide valuable labor. Ems with different capabilities are assigned to different tasks; groups of ems are organized into functional teams; cities are built to collocate ems for enhanced productivity. Ems are paid for their services and can spend their earnings on leisure, though competition between ems drives wages down.

While these basics are broadly similar to much of today's society, the nature of ems creates some radical differences. For starters, while some ems are housed in robotic bodies to perform physical-world tasks, most live and work disembodied in virtual reality worlds. This enables them to travel at great speeds via telecommunications lines. Additionally, ems can be sped up or slowed down as needed, causing them to experience the world (and act on it) thousands or millions of times faster or slower than biological humans. Finally, ems can be copied repeatedly, which enables em teams to perform extreme multitasking. The book even proposes the concept of "spurs", which are em copies that perform a brief task and are then terminated or retired.

These novel capacities enable ems to vastly outperform biological human workers. Em labor thus displaces humans while creating rapid economic growth. The book proposes that those humans who own parts of the em economy "could retain substantial wealth", but everyone else is "likely to starve" (p.336). To my eyes, this seems overly optimistic: if ems are so much smarter than humans, and if they have such a profit motive, then surely they could figure out how to trick, threaten, or otherwise entice humans into giving up their wealth. Humans would likely die out or, at best, scrape by in whatever meager existence the ems leave them with.

Given these dire prospects for humans, one might question whether it would be good to create ems in the first place. Unfortunately, the book does not consider this topic in any detail. The book is pro-em, even proposing to "subsidize the development of related technologies to speed the arrival of this transition [to the em era], and to subsidize its smoothness, equality, or transparency to reduce disruptions and inequalities in that transition" (p.375). But this position is tenuous at best and quite possibly dangerous.

The book's methodology is notable for its emphasis on social science. Many people who write about topics like uploading are trained in physical science, computer science, or philosophy; they thus "don't appreciate that social scientists do in fact know many useful things" (p.34). Hanson's defense of social science in this context is welcome.

The book's methodology is rooted in extrapolation of current social science to em world conditions. At times, the extrapolations seem strained. For example, a claim that economic benefits of urban agglomeration would continue in an em world (p.215) cites Morgan (2014),

which is a popular media description of Amazon's large data centers. It is true that Amazon data centers are large, and it may well be true that em cities are large for similar reasons, but the latter does not necessarily follow from the former. One can sympathize with Hanson, who is trying to present the social science of a topic that virtually no social scientists have previously touched. Indeed, the breadth of research covered is impressive; the book quietly doubles as a good literature review. However, some of the language relating the literature to the em world could have been better finessed.

In other stretches, Hanson's personal tastes are apparent. This is seen, for example, in discussions of combinatorial auctions and prediction markets (p.184-188), two schemes for adapting market mechanisms for social decision making. Prediction markets in particular are a longstanding interest of Hanson's. The book's discussion of these topics says little about the em world and seems mainly oriented towards promoting them for society today. The reader gets the impression that Hanson wishes society today was more economically efficient and rational in a certain sense, and that he has embedded his hopes for a better world into his vision of ems.

In conclusion, *The Age of Em* is a thoughtful and original work of futures studies, bringing an insightful social science perspective to the topic of mind uploading. The book covers wide ground about the nature of the em world, offering a valuable resource for anyone interested in the topic, or about futures studies more generally. The book is accessibly written and could be read by undergraduates or even advanced high school students, though experts from any discipline will also find much of interest. The book is especially worthwhile for anyone who could benefit from an overview of contemporary social science research, which the book provides in abundance.

The book is not flawless. One can disagree with some of the book's arguments, including some of the important ones. The author's personal perspective is apparent; different authors would often reach different conclusions about the same topic. The organization is often scattered, and some literature is misused. However, none of these flaws take away from the significant accomplishment of describing the em world in great detail. The book offers a starting point—indeed, a baseline—for the serious study of future uploading scenarios. One can only hope that additional scholars, especially in the social sciences, follow the book's lead in taking on this important future scenario.

Acknowledgments

Tony Barrett, Roman Yampolskiy, Ted Fuller, and an anonymous reviewer provided helpful comments on an earlier draft. Funding was provided in part from Future of Life Institute grant 2015-143911. The views in this paper are the author's alone and not necessarily those of the Global Catastrophic Risk Institute or the Future of Life Institute.

References

- Morgan, T. P. (2014). A rare peek into the massive scale of AWS. *EnterpriseTech*, November 14.
- Schoemaker, P. J. H. (1995). Scenario planning: A tool for strategic thinking. *Sloan Management Review*, 36(2), 25-40.
- Steinhardt, J., & Shapiro, M. A. (2015). Framing effects in narrative and non-narrative risk messages. *Risk Analysis*, 35(8), 1423-1436.