

Robin Hanson (2014) What Will It Be Like To Be An Emulation? In Russell Blackford, Damien Broderick, ed. *Intelligence Unbound: The Future of Uploaded and Machine Minds*, Wiley.

What will it be like to be an Emulation?

Robin Hanson

The next big epochal change is likely to be the arrival of “artificial intelligence”: machines intelligent enough to substitute wholesale for human workers. A good guess about the origin of such machines is that the first ones, arriving within roughly a century, will be uploaded humans—whole brain emulations (Hanson 1994, Hanson 2008; Sandberg and Bostrom 2008), or “ems.”

Definition: An *em* results from taking a particular human brain, scanning it to record its particular cell features and connections, and then building a computer model that processes signals according to those same features and connections. A good enough em has roughly the same input-output signal behavior as the original human. One might talk with it, and convince it to do useful jobs.

Rather than debate the feasibility or timing of ems, or ponder their implications for the philosophy of mind or identity, this chapter will seek realistic social implications. In what sort of new social world might ems live, and what will it be like to *be* an em?

Why think about ems now?

We take far more effort to study the past than the future, even though we can’t do anything about the past. But modest efforts have often given substantial insights into our future, and we would know much more about the future if we tried harder. If policy matters, then the future matters, because most policy effects happen in the future. And unless we are ridiculously self-centered, the distant future matters the most, since with continuing growth we expect that the vast majority will live there.

If, within decades, enormous changes will happen, policy analyses that ignore such shifts might be irrelevant or badly misdirected. So it’s important to try to foresee any big upcoming changes, and their likely consequences. My analysis takes the somewhat unusual approach of using basic social theory, in addition to common sense and trend projection, to forecast future societies.

Understanding our descendants also helps us to understand our place in history. We define ourselves in part by how we differ from our neighbors in time.

Life among the ems

For some portion of the next century, the world still looks much like ours. Then a few firms start making cheap mind-emulating machines able to substitute wholesale for human workers. States or nations that do not prohibit those machines quickly get fantastically rich, and then dominate the world economy.

Billions of ems, perhaps trillions, are now found in cities packed dense with computer

hardware. Such cities are very hot, even glowing, covered and pressurized, and arranged as a vast three-dimensional lattice, rather than as separate buildings, pulling winds of hot air into tall clouds overhead.

Em cities are likely toxic to ordinary humans, who, controlling most of the rest of the Earth, mostly live comfortably on their em-economy investments. While ems could easily buy the rest of the Earth, they don't care enough to bother, beyond ensuring energy, raw materials and cooling for their cities.

While some ems work in tiny robotic bodies, most do so in virtual offices, and almost everyone plays in virtual reality. Here is the surprising thing: either way, ems think and feel like humans. Their world looks and feels to them much as ours looks and feels to us. Like humans, ems remember a past, are aware of a present, and anticipate a future. Ems can be happy or sad, eager or tired, fearful or hopeful, proud or shamed, creative or derivative, compassionate or cold. Ems have friends, lovers, bosses, and colleagues. Em mental and psychological features might differ from the human average, but, despite some eccentrics, are usually near the range of human variation.

Ems reproduce by making copies who remember exactly the same past, but then diverge afterward with differing experience. Typically, whole teams of copies will work and socialize together, and then retire together from the workforce. Often, hundreds of similar teams are made at once. Most ems remember formerly being in another team, and agreeing to let this copy be made so he or she can join the new team. Ems feel grateful to exist, and mostly accept their place in the world.

As with us, most em choices are Either/Or: shall I spend my personal budget of time and money to do this, or that? An em might pick an entertainment to watch (or participate in), or a friend to be with. But for ems, bigger life choices can be Add/And: shall I allow another me to have a new life? Big choices are often made jointly with copies, teammates, and investors. Teams are tightly bound, like traditional working class work groups, with lots of insults, teasing, and profanity to express and test emotional strength.

On the upside, most ems have office jobs, and work and play in spectacular high-quality virtual realities without hunger, severe cold or heat, grime, physical illness, or intense pain—unless those are chosen deliberately. Unlimited lifespans are possible. On the downside, em wages are so low, for reasons we'll consider in a moment, that most barely scrape by while working hard for half or more of their waking hours (assuming that ems still need sleep and dreaming time). And their extended lifespans induce greater wealth inequality among ems.

After a short time, most ems, due to strong selection pressures, are copies of the thousand or fewer humans best suited for em jobs. These thousand-or-so copy “clans” are the smartest, most cooperative, hardest-working people comfortable with often splitting off a short-term “spur” copy to do a several-hour task and then end (die), or maybe retire to a much slower speed. At any one time, most working ems (weighing by count or speed) are spurs.

The culture of ems

An em would trust its clan of copies more than other groups and may get clan-based life coaching, drawn from the experiences of millions of similar copies. Since clans might be legally liable for member actions, they regulate member behavior to protect clan reputation, making ems pretty trustworthy. To avoid nepotism, inefficiency, and social weirdness, most workplaces would avoid having ems from the same clan take on many different roles. Em firms would be

bigger, with more managers, more meetings, and narrower job descriptions.

Em minds run at many different speeds, with faster giving higher status; different speeds have different cultures, and fast cultures fragment into local cultures. The cost to run an em is nearly proportional to its speed over a wide range, making the cost of a subjective minute of em experience nearly independent of speed. Because of different speeds, one-em one-vote won't work, though speed-weighted voting might do. The politics of em city-states and firms might reflect shifting coalitions of clans, perhaps coordinated by widely watched one-copy-per-clan "reality show" social gatherings.

To allow romantic relations when there is unequal demand for male vs. female workers, the less demanded gender might run slower, and periodically speed up to meet with faster mates. Em sex would be only for recreation; ems would have superb bodies, and most would have impressive minds. But few would have the youthful personality or wealth that men and women today find sexually attractive.

The em economy might double every month or week, a growth driven less by innovation, and more by population growth. While this em era might only last for a year or two, for a typical em running a thousand times as quickly as a human, it would last for subjective millennia. To such ems, their world seems more stable to them than our world seems to us. Bosses and software engineers run even faster, some at a million times human speed.

Em industries include security, emergency-response, training, law, finance, news, entertainment, politics, education, software, computer and communication hardware, energy, cooling, transportation, manufacturing, construction, and mining. Em tasks, to the extent that these have not already been taken over by smart but not human-level AIs, include designing, marketing, selling, purchasing, managing, negotiating, administration, testing, monitoring, diagnosis, repair, cleaning, driving, accounting, assembling, packing, installation, mixing, sorting, and fitting (to the extent that these tasks have not been totally automated).

Faster ems with physical but non-biological bodies have proportionally smaller embodiment; a typical em operating a thousand times quicker than a human would stand two millimeters tall. Physical ems see a mix of reality and virtual overlays. Most long distance physical travel is "beam me up" electronic travel, though done with care to avoid mind theft. Compared to humans, ems are far less fearful of the death of their particular copy. But the risk of mind theft is a frightful evil, being a plausible route to personal destitution, torture, or worse, and a serious threat to the economic order. While some ems offer themselves as open source and free to copy, most ems work hard to prevent or mitigate mind theft, a cause that evokes great moral feelings and charity efforts.

Humans today reach peak productivity around the age of fifty. Most ems are near their age of peak productivity, perhaps sixty or more. Older minds become less flexible, and must eventually retire to an indefinite life of leisure at a much slower speed, sometimes as low as one thousandth of human speed. The subjective lifespans of both humans and slow em retirees depend mainly on the stability of the em civilization; a collapse or revolution might kill them. Ems would be wary of allowing theft from retirees and humans, as that might threaten the institutions ems use to keep peace among themselves.

Em identity

Ems focus their identity less on individual personalities and abilities, and more on being part of a particular team. While copy clans coordinate to show off common clan abilities, individual ems

focus on showing their identity, abilities, and loyalties as team players. Teams prefer to socialize internally, to reduce team productivity variance. Instead of trying to cure a depressed or lovesick em, that individual can be reverted to a version from before the problem appeared. For larger problems, whole teams might be reverted or replaced.

Team members can often read the surface of each other's minds. In general, ems specialize in dealing either with high transparency or high opacity of associates' minds. Ems can mistrust their own experiences, suspecting that unusual experiences might be simulations designed to test their loyalty or extract a secret. Many schemes are tried to avoid being fooled by a simpler automated mimic.

By having one em plan and train, and then split into many copies who do the related tasks, ems find it easier to prepare for and coordinate tasks. For example, a single em can conceive of a software or artistic design and vision, and then split into an army who execute that vision. Childhood and job training are similarly cheaper; copies of young humans' minds are rare but pampered.

Larger projects can also be completed more often on time, if not on budget, by speeding up the ems who work on lagging parts. More generally, em firms are larger and better coordinated, because fast bosses can coordinate more, and because employee clans with strong financial and reputational interests less often game employee evaluations.

Even if most ems work hard much of the time, and will probably end or retire soon, they usually remember a long history of succeeding against the odds, with long stretches of leisure and training punctuated by short intense work periods. Such self-deception is easier for ems because when things really matter spurs can convince them to do what is practical without explaining why. To most ems, it seems good to be an em.

Inequality

How unequal would ems be from one another?

During humankind's forager era, the main units of organization were bands of roughly twenty to fifty people, and smaller family units. Foragers did differ in personal property or prestige, but lacking larger units foragers did not have inequality analogous to our unequal towns, firms, or nations. Over the million plus year forager era, however, foragers had great inequality of lineages, in the sense that almost all lineages went extinct.

The farmer era saw larger units of organization, such as clans, towns, nations, and empires. While empires sometimes became nearly as large as feasible, transportation limits weakened their influence on local behavior. Towns were usually much smaller than nations, and firms were smaller still. While most farmers lived in rather small towns, with the emergence of industry people aggregated evenly across feasible city sizes.

For most products today, market shares are relatively concentrated within transport-cost-limited market areas; for each type of product, a small number of firms supply most customers.

Power laws often well describe the upper tails of the distributions of how such items are grouped into units. In such cases, a power of one describes a uniform distribution of items across feasible unit sizes. Powers greater than one describe more equal distributions, where most items reside in small units, and powers less than one describe less equal distributions, where most items are clumped into fewer larger units. Compared to some reference-sized unit, for a power of one a unit twice the size is half as frequent. For a power greater than one, such double size units are less frequent, and with a power less than one they are more frequent.

In all eras so far, family names have been distributed relatively equally, with a power of about 2.0. In the late farmer era, towns were distributed with a power of about 1.4 (Nitsch 2005), roughly the power that describes today's individual wealth distribution (Davies et. al. 2008). Today, nations, firms within nations, and cities within nations, are all distributed with a power of about 1.0 (Axtell 2001; Eeckhout 2004; Giesen 2010). Product suppliers are usually more clumped, with a power less than one.

Em-era clans probably clump much more unequally than did prior-era families; most ems, as noted, might come from a dozen to a thousand clans. Labor markets become much more like today's product markets, with a few clans supplying most workers in any given skill area. Clans coordinate to prevent the theft of their members, and to ensure that they can profit from training investments. Fixed costs of training and marketing create concentration and market power, though not net profits after paying for fixed costs. Because virtual reality meetings greatly cut the commuting congestion that limits city sizes today, em cities and nations are also likely to be much more unequal, with most ems found in a few very large city-states.

Em firms would probably continue to be distributed evenly across feasible firm sizes. Even so, because a larger economy could support larger firms, typical firm size would increase. Today in a world of seven billion people, firms are spread evenly between, roughly, one and one million employees. The middle firm size, where half of workers are in larger firms, is about a thousand. For every factor of one hundred increase in the size of the largest feasible firm, there should be about a factor of ten increase in the size of the average size firm.

Em wealth would also probably be more unequal, because of both indefinite em lifespans and the em capacity to run at different speeds. Today, cities and firms can spread uniformly across all feasible sizes due to their indefinite lifespans; if a city or firm continues to succeed, it can grow, regardless of its age. In contrast, while wealthy individuals today often consistently grow their wealth over their lifespan, their children usually fail to continue the same growth pattern. Rags-to-riches-to-rags in three generations is, after all, a common scenario. With indefinite em lifespans, however, successful ems could continue their winning financial habits indefinitely. This should tend to distribute wealth more evenly across all feasible wealth sizes, rather as firms and cities are distributed evenly today across feasible sizes.

The tendency of em minds to become inflexible with experience, however, could somewhat reduce this effect, especially if rich ems do not reduce their speed. On the other hand, the ability to run at many different speeds will expand the possible range of em wealth. While very poor humans must starve, very poor ems can run at very slow speeds or be archived at negligible cost.

A power of one for the distribution of em wealth would make wealth more unequal than today, both because this is a more unequal power, and because a richer em world can make feasible both higher and lower wealth levels. While today's richest person holds ~0.02% of the world's wealth, the richest em may hold a much larger fraction, perhaps ~ 2%. The exact fraction, however, might depend on em wealth sharing arrangements. For example, some em clans might agree to share much of their wealth internally.

War

Today the frequency of war, civil war, and criminal violence tends to fall as nations get richer, have older citizens, and sit more toward the extremes of democracy and autocracy (Magee and Massoud 2011). The fact that most ems would be near peak-productivity subjective age suggests

that their attitudes toward war would be more like that of a typical fifty year old, which today is less supportive of war. The fact that ems would be poorer in many ways, however, might push for more war. Other relevant factors might include em gender imbalances, reduced calming effects from raising children, productivity tweaks that might increase aggressiveness, and a generally more competitive world.

Having ems concentrated in a few very large dense cities should reduce the frequency of small wars. After all, wars within cities are very destructive of the value that cities produce, and cities have many other ways to help discourage and settle internal disputes. Wars were especially common in the farming era because valuable loot could be seized, and because farm land quickly returned to full productivity even after all-out war burned crops and buildings to the ground.

Wars between big cities are a greater concern, in part because some cities might beat others in economic competitions. Nuclear weapons could still threaten to destroy entire cities in a single strike. To a kilo-em, however, running at a thousand-fold speed-up, a nuclear missile that gave fifteen minutes of objective warning time would give ten days of subjective warning, making it much easier to respond flexibly. In contrast, lasers and directed energy weapons would still have apparently instantaneous effects on city scale distances. To mega-ems, however, such light speed weapons would seem to have substantial delays on such scales.

Em soldiers would hardly fear the death of any one copy, as they could easily revert to a recent backup copy. The material resources lost when a copy (or a city) is destroyed, however, would still sting. Wars of attrition would be accounted less by deaths than by resources destroyed.

Clans

How would clans organize?

Each of us today is part of many organizations: neighborhoods, firms, clubs. But we rely most on our families when we seek strong long-term bonds and trust. It is in families that we mainly share resources, let ourselves be most vulnerable, and seek help in bad times. Humans evolved to trust families more than others because of closer genetic relations, and we have developed many family-specific adaptations to complement such trust.

Today, identical twins are even more closely related than other family members. Because of the rarity of such twins, humans have probably evolved few adaptations specific to twinning. Even so, the bonds and trust between identical twins are usually even stronger than between other siblings and parents.

Ems will have access to a new and tighter unit of organization: clans of copies of the same original person. Compared to families or even identical twins, ems will have yet stronger reasons to trust and bond with fellow clan members. This makes em copy clans a natural unit of finance, reproduction, legal liability, and political representation.

Clans with billions or more members would likely split into sub-clans of millions or thousands, at least for many purposes. The tree structure of their copy ancestry is a natural basis for such sub-clan grouping. During the farming era, ancestry trees of family clans were often a basis for legal liability, political coalitions, and much else. Relative to em clans, em sub-clans would be more similar to each other in jobs, hobbies, friends, personality, and shared memories.

Law usually tries to deter crime via small chances of big punishments. It can be especially hard, however, to impose big punishments cheaply on short-lived ems (such as spurs). Giving spurs a longer, fast life to suffer judicially-imposed retributive torture, for example, could

be expensive. For such ems it can make sense to use vicarious liability, i.e., to have legal liability apply to an associated unit with deeper pockets. In principle, this could be any sponsor willing to take on the liability, but in practice it probably would be the sub-clan containing this em. The em would care more about harming its sub-clan, and the sub-clan would know better how to discourage liability-inducing events.

Holding closely related copies responsible might seem more legitimate if archived versions could be tested in simulation to see if they would behave similarly given a similar criminal opportunity, at least for situations where behavioral inclinations seem relevant. Em clan liability is similar to the way that farmer societies often made larger family units liable for the crimes of family members. In response, clans legally liable for member behavior would seek to regulate problematic behaviors, and seek sufficient powers in order to manage clan reputations. Ems in clans that more strongly manage their reputations are likely to be more predictable and reliable.

Since political regimes become unstable when members do not accept their legitimacy, and since we tend to accept and defer to internal decisions of families, more than to other kinds of groupings, farmer-era city and regional politics often used families as the first unit of grouping. Accordingly, ancient politics was often the politics of shifting coalitions of family clans (Braekvelt 2012). This remains true in the Middle East and Asia (Sailer 2003).

For ems, it makes more sense to use clans as the first unit of political organization. If most ems constitute from a dozen to a thousand clans, the politics of an em city-state might naturally be the politics of shifting coalitions of em copy clans. Clans whose root ems were raised together as children might tend to ally more with each other, and childhood cohorts might be chosen with this prospect in mind. To encourage clan loyalty, clans could also try to anticipate and coordinate to prevent situations where clan members compete over friends, lovers, or jobs. For example, different sub-clans might be prevented from competing for the same jobs or teams. In a clan-dominated era, people whose basic personality makes it harder to coordinate and compromise with copies of themselves might suffer a serious competitive disadvantage.

Intelligence unbound

Merely being faster will make ems seem smarter in many ways. Even if we control for speed, however, ems would be smarter than humans because of stronger selection, both in choosing ems from among humans, and in selecting mental tweaks for the emulation process. Ems might also find ways to get smarter by expanding their mental hardware, perhaps by increasing the number or size of many simple repeating brain circuits. And some ems might choose to lower their brain hardware costs by getting dumber, via reducing the number or size of such circuitry.

We expect smarter workers to accomplish more with the same resources, to make fewer and less consequential mistakes, to master a wider scope of tasks and skills, to communicate effectively with a wider range of fellow workers doing different jobs, to adapt to changing circumstances faster and better, and to learn specific and specialized roles faster and better.

How these changing capacities would change the mix of jobs and careers will vary by profession and industry. In more stable, slower-changing environments, the gains from making fewer mistakes and better learning specific skills should allow a finer division of labor into more specialized, more interdependent roles. By contrast, in uncertain and rapidly changing environments, the abilities to master more skills and adapt faster should matter more; em organizations should be able to function effectively in a wider range of such environments, using

smaller teams of less specialized workers.

By being smarter, ems should become better at innovation, em bosses should be able to manage simultaneously a wider range of subordinates, em jobs could be redesigned more often, and em careers should last longer before em minds become too fragile to compete effectively.

But note: while these changes are non-trivial, they are also not radical. A society full of creatures much smarter than we are could remain quite recognizable and understandable to humans. It need not create a horizon of visibility beyond which we cannot foresee. Some claim that a world of much smarter creatures would quickly become a world run by a single dominant mind. I've explained my reasons for that elsewhere (Hanson and Yudkowsky 2013), but even if it happened, such a world would differ primarily in its concentration of power, not in the intelligence of its creatures.

Conclusion

The analysis in this chapter suggests that lives in the next great era might be as different from what we are accustomed to as ours are from farmers' lives, or farmers' from foragers'. Some readers, living industrial era lives and sharing industrial era values, will be disturbed to see a forecast of em era descendants with choices and life styles that appear to reject many values they hold dear. Such readers might be tempted reject this scenario, and fight to prevent it, perhaps preferring a continuation of the industrial era to the arrival of such a different epoch.

But I advise such readers to first try hard to see this new era in some detail from the point of view of its typical residents. See what they enjoy and find meaningful, consider what fills them with pride, and listen to their criticisms of your era and values. Understand how disturbed foragers were by many aspects of farmers' lives, and farmers by many aspects of our industry era lives. And then decide.

References

Axtell, Robert L. 2001. Zipf distribution of US firm sizes. *Science*, 293, p. 5536

Braekevelt, Jonas, Buylaert, Frederik, Dumolyn, Jan, and Haemers, Jelle. 2012. The politics of factional conflict in late medieval Flanders. *Historical Research*, 85(227), pp.13–31.

Davies, James B., Sandström, Susanna, Shorrocks, Anthony, and Wolff, Edward N. 2008. The world distribution of household wealth, UNU-WIDER Discussion Paper 2008/03, February.

Eeckhout, 2004. Gibrat's law for (all) cities. *American Economic Review*, January, pp. 1429–1451.

Hanson, Robin. 1994. If uploads come first. *Extropy* 6(2), pp. 10–15.

Hanson, Robin. 2008. Economics of the Singularity, *IEEE Spectrum*, June, pp. 37–42.

Hanson, Robin, Eliezer Yudkowsky. 2013. The Hanson-Yudkowsky AI-Foom Debate. Berkeley, CA: Machine Intelligence Research Institute.

Magee, Christopher and Massoud, Tansa. 2011. Openness and internal conflict. *Journal of Peace Research*, 48, pp. 59–72.

Nitsch, Volker. 2005. Zipf zipped. *Journal of Urban Economics*, 57(1), pp. 86 –100.

Sailer, Steve. 2003. Cousin marriage conundrum: The ancient practice discourages democratic nation-building. *The American Conservative*, January 13, pp. 20–22.

Sandberg, Anders and Nick Bostrom. 2008. *Whole Brain Emulation: A Roadmap*, Technical Report #2008-3, Future of Humanity Institute, Oxford University, www.fhi.ox.ac.uk/brain-emulation-roadmap-report.pdf (accessed October 26, 2013).

Biographical Note

Robin D. Hanson is an associate professor of economics at George Mason University and a research associate at the Future of Humanity Institute of Oxford University. He is known as an expert on prediction markets and was a principal architect of the Foresight Exchange, DARPA's FutureMAP project, and IARPA's DAGGRE project.