

SILICON SIMULACRA

Post-humans of the Machine Worlds

Len Ellis

Conclusion

Copyright 2010 by Len Ellis

All rights reserved.

Subject to the exception immediately following, this book may not be reproduced, in whole or in part, including illustrations, in any form (beyond that copying permitted by Sections 107 and 108 of the U.S. Copyright Law and except by reviewers for the public press), without written permission from the publishers. The author has made available online versions of the book's chapters under a Creative Commons Attribution Noncommercial-No Derivative Works 3.0 United States License.

20 Charles F. Manski, “Interpreting the Predictions of Prediction Markets,” August 2005, www.aeaweb.org/annual_mtg_papers/2006/0106_1015_0703.pdf and Justin Wolfers and Eric Zitzewitz, “Interpreting Prediction Market Prices as Probabilities,” January 8, 2007, <http://bpp.wharton.upenn.edu/jwolfers/Papers/InterpretingPredictionMarketPrices.pdf> provide mathematical assessments of using trading prices as a stand-in for probabilities.

21 Jennifer Watkins, “Prediction Markets as an Aggregation Mechanism for Collective Intelligence,” (speech, 2007 UCLA Lake Arrowhead Human Complex Systems Conference, April 25–29, 2007) <http://repositories.cdlib.org/hcs/WorkingPapers2/JHW2007> is a particularly clear exposition of these requirements. James Surowiecki, *The Wisdom of Crowds* (New York: Anchor, 2005) emphasizes how hard and rare it is to get everything right.

22 To be accurate, the commodity exchanges enable companies to hedge against price changes in raw materials as well as enable traders to speculate.

23 The use of money as an incentive is not, however, necessary for prediction markets to generate accurate results; see Emile Servan-Schreiber et al., “Prediction Markets: Does Money Matter,” *Electronic Markets* 14 (September 2004).

CONCLUSION

What is appearance for me now? Certainly not the opposite of some essence: what could I say about any essence except to name the attributes of its appearance?

—Friedrich Nietzsche, *The Gay Science* (1882)

In 1993 the *New Yorker* magazine ran a cartoon by Peter Steiner in which one dog, sitting in front of a home computer, says to another dog sitting beside him, “On the Internet nobody knows you’re a dog.” The caption alluded to the chat rooms and electronic bulletin boards of the text-only pre-Web Internet where users, logged on under their screen names, discussed topics of all sorts. Among the obvious benefits, then and now, ordinary folks in socially stigmatized situations, such as living with diabetes or filing for bankruptcy, can get or give something helpful without disclosing their real world identities. Even more broadly, ordinary folks in ordinary situations, such as shopping for a car or finding new dinner recipes, can get what they want without dealing with salespeople or being obliged to reciprocate.¹ On the downside, the use of screen names enables sexual predators, con artists and other deceivers to pretend to be someone they are not.

Internet users are largely lackadaisical about our online anonymity. Losing it would be a disaster, but few pay it attention. Large majorities tell pollsters that they are “concerned” or “very concerned” about privacy online. In actuality, however, vast numbers of us give up our birthdays, zip codes and other personally identifying data to any Web site that asks in exchange for ringtones, screen savers, horoscopes and other trinkets. Similarly, very few use the privacy-protecting tools that have long been available. Leaking out personal data has become just a routine feature of the contemporary consumer condition.

We have a greater stake in anonymity's opposite, authentication. Verifying one's identity is essential to every electronic transaction, offline and online, and the recent spread of terrorist attacks has made it essential to our physical safety as well. The obvious downside of authentication systems is identity theft. The increasing importance of verifiable identities may help in perceiving the emergence of our silicon simulacra, but they are very different. Authentication verifies us. Our simulations make us knowable; they turn us out in public to others, albeit inside machines.

I am a number!

Authentication dates back to the 1920s. It was one dimension of the in-house credit systems that department stores offered their good customers. Some people at the time were nervous about the gathering of personal information for these services, but convenience won out. Authentication became ubiquitous in the 1970s as plastic credit cards replaced cash in everyday purchases. We still don't know the details. We just know that every swipe sends the card number somewhere where it gets processed and then sent back, either approved or denied.² Similarly, when we register at Web sites, we know, sort of, that the valid e-mail address we provide is an individual identifier and that the password we provide verifies it. We resonate to the defiant protest, "I am not a number" as a cultural critique, but we accept that the great circuits know us as numbers, two numbers in particular.

The first is one's Social Security number or its equivalent, like the number of one's passport, military service ID or state-issued driver's license. Because these numbers are unique and assigned to individuals, they can verify that a person is who she claims to be and can function as keys to open and access accounts and dossiers. The other is one's credit score. Compiled by three consumer credit reporting agencies—Equifax, Experian and Transunion, this number

is a rating; it compares each of us to others and on that basis grants access to other resources, usually more or less credit. Not long ago, a person's credit score was inaccessible. Today, consumers have a legal right to their credit reports for free at least once a year and can challenge and correct the information they contain. For a fee, consumers can also subscribe to services that will monitor and report on a daily, weekly or monthly basis which third parties are requesting one's credit score; one can even place certain access controls on such requests, including an outright freeze.

As authentication regimes grew prevalent, so did identity theft. We tolerate this downside but shouldn't. This crime *du jour* is not an inescapable consequence; rather, it reflects how informational rewards and risks are defined. The problem that gives rise to identity theft is the separation of rewards and risks: businesses get the rewards while consumers bear the risks.

Legally, no individual can be a victim of identity theft. Theft applies to property, and under U.S. law the individual has no property interest in her data. The companies that compile the data into databases have the property interest. In the landmark legal case, American Express was exonerated for selling its cardmembers' names to merchants because, the court argued, "an individual name has value only when associated with the defendant's lists." That is, by compiling and categorizing our personal information, the data aggregator creates the value that constitutes the property interest.

But the owners of this property have few incentives to protect it. Although most states require companies to report data breaches and threaten to impose fines on those who don't, there are many holes. Some states require only businesses in certain industries to report while still other states require it only if the business suspects the stolen data will be used to commit fraud. What's more, even when fines are levied for failure to report a breach, they're small, much less than what it would cost to research and report a breach. So, many data thefts go unreported; it's cheaper to just pay the

fine. In short, the downsides for property owners are minor if their property is stolen.

In contrast, the consumer becomes a victim when someone uses the stolen data to impersonate the consumer. That crime is not theft but fraud, and the harm is relatively major, usually charges against one's credit cards or, worse, entirely new credit card accounts and bank loans opened in one's name. In addition, significant time is required to repair the informational damage, that is, to correct one's records at banks, retailers, credit rating agencies and police departments. What's more, seeking compensation from the company whose security was breached is not viable because connecting the fraud back to the theft is difficult, usually impossible.

This disconnect—businesses get the rewards while consumers get the risks—encourages both theft from the former and fraud against the latter. Rather than address the cause, we're letting the marketplace address the effects by asking prospective victims to pay in advance for protection. Quite concretely, consumer services for identity protection and credit monitoring are flourishing businesses.

While Americans have long been comfortable in being individually verified as consumers, the terrorist attacks of September 11, 2001 have made authentication necessary for us as citizens as well. The age of terror has made anonymity problematic and intensified a classic dilemma: government surveillance to ensure our collective security encroaches upon individual privacy. Immediately after 9/11, for example, the Defense Department wanted to datamine all the transaction records of everyone in the country, looking for tell-tale patterns of terrorist plotting, while the U.S. Attorney General wanted to allow the FBI not only to conduct surveillance outside specific investigations but also to compile dossiers about individuals who participate in public assemblies and who visit certain Web sites.

The government's track record isn't good here. Back in the 1950s, '60s and into the '70s, the FBI used its surveillance

capabilities to undermine legitimate political activity, including both the civil rights and anti-war movements. But no government's track record was or ever will be good here. The peril is inherent and the Founding Fathers knew it. From the grievances listed in the Declaration of Independence to the checks and balances built into the Constitution to the individual freedoms secured by the Bill of Rights, all our charter documents focus on protecting us against the ever-present temptation of governments to encroach upon the rights of those they are supposed to protect.

Whenever the watchdogs bark, we should listen, but our need for increased security must also be addressed. Defending individual privacy against government surveillance is necessary; it keeps the tension healthy, but it does not solve the dilemma. Nor is the solution to demand anonymity in public life. Freedom means little if its exercise requires anonymity. Rather, we have freedoms of speech and assembly to the extent that we can stand up and be counted without fear of sanctions.

Not the right to privacy but the duty of publicity applies here. The American jurist Louis Brandeis coined both phrases. The latter means that people have a right to know who is addressing them and that speakers have a duty to be accountable. The sunlight of others' approbation exercises a disciplining pressure. Anonymity removes it. As many know from online discussion forums, anonymity facilitates all sorts of disruptive and destructive behaviors. From our experience of computer systems generally, many also know the opposite, presenting a valid ID grants access to whatever resources one is entitled. "Enrollment enables entitlement" is a well-known motto in computer security circles. In the wake of 9/11, we've all grown accustomed not only to more surveillance but also to presenting some form of official identification whenever we try to cross some threshold, online and offline, and gain access to whatever's on the other side. In balancing safety and privacy, our willingness to be individually authenticated is emerging as the way forward.

The rise of identity theft and our response to terrorism have made us aware that each of us has identities that machines can read. These identifiers make us safer by authenticating us; at the same time they render us vulnerable to impersonation and fraud. Identity is not individuality, however, and authenticating the one does not disclose the other. Under authentication regimes, we remain private individuals. Our silicon simulacra, the data profile and the cyberpersona, are different. They make us visible and knowable as individuals to others.

The New Publicity

In general, Americans don't go in much for public life. We don't have the pub cultures of northern Europe or the sidewalk social life of cafes and bistros common in the olive zone. As for our home-grown civic associations, fraternal orders, business clubs, ladies' auxiliaries and other formally organized groups, they've long since faded away. Americans may have once been a nation of joiners, but today we bowl alone. Indeed, it is a common complaint that Americans lead largely privatized lives.

Both the datascape and cyberspace, however, want us knowable. Although the former relies largely on surveillance while the latter promotes self-disclosure, the end result is the same. Both subvert the essence of the Romantic self—mysterious, quixotic and out of reach—and replace it with a self that is knowable, present in the here and now and open to entrance by others.

In the datascape such visibility enables those with power—governments, businesses and employers—to make ever-more precise decisions that benefit themselves rather than citizens, consumers and employees. Cyberspace is similar. Every Web 2.0 application invites each of us to make our activities, interests, tastes, preferences, expertise and predictions available and knowable to others, including marketers who buy our eyeballs and foot the bill for these free applications.

But the ways in which we become knowable inside each machine differ and not in the ways many would expect.

Contrary to the concerns of privacy champions, the datascape does not threaten individuals. Individuals with our qualities and possibilities don't even get into databases; we're left outside. Only machine-readable attributes make their way in, and, once inside, they just sit idle as raw material until someone queries the database. Only then is the raw material processed. Those real-life people whose attributes conform to the query become the hypothetical persons of the data profile. This is a probabilistic informational entity: a set of persons who are statistically predicted to be more likely than others to conform to some specific end state desired by some marketer, bureaucrat or other planner. On that basis, it provides the scientific (read: objective and impartial) rationale for those decision makers to treat some of us differently than others. To be sure, statistical differentiation for commercial discrimination affects us individuals. Some of us get treated better than others, but that happens to us as members of groups. Individuals are not at risk. Rather, the direct impacts of this decision-making regime are largely social. Three are especially problematic.

First and foremost, while using data to understand and govern human affairs accomplishes a world of good, it does not capture our qualities or our possibilities. The omission has two consequences. Ersatz versions of us go into databases, and the ersatz answers about us that come out are the basis for marketers, bureaucrats and other planners to make decisions about nearly every inch of the built environment in which we live. The built environment produced by their decisions fits the data well enough but doesn't quite fit the humans from whom the data was abstracted. The second consequence of this omission is worse. The more data is used to apprehend and administer human affairs, the more it tends to push into the shadows other traditions that do address our qualities and possibilities. This surreptitious coercion into cultural amnesia could be characterized

as ideological because the hallmark of any ideology is to prescind alternative ways of apprehending the world. Ideological or not, this systematic forgetting disadvantages us in trying to make sense or meaning of our lives.

The second social problem derives from the one overarching purpose of all data-based decision making, to create and measure difference. Differences are not natural. Specific decision makers hypothesize specific differences for their specific purposes; they then create variables that ostensibly express those differences, collect observations and analyze the results. Whatever we have in common is of no use to them because it cannot support a decision to allocate resources here *rather than* there. So they never look for what we share or may want to share. Only difference is hypothesized, only difference comes out and only difference is made real via the data-based decisions about the built environment. Whatever the specific differences hypothesized in any particular instance and whatever one thinks of those differences, a society centered on the cultivation of difference is going to have a hard time figuring out how to move forward together, at home or abroad.

The third socially problematic aspect is the rarely discussed vision of the larger society within which difference is framed. That vision sees society as a socioeconomic stratification system in which the established hierarchy of power and interests are givens. Although this framework suits the purposes of those who own, operate and utilize the apparatus, it's only one way of seeing our society. Moreover, it omits several socioeconomic changes of historical magnitude that actually challenge us today and, like all feedback systems, is self-reinforcing and self-perpetuating. In short, the framework locks us into the past and tends to preempt different visions of the social fabric.

All three consequences—the cultural amnesia, the cultivation of difference and the outdated vision of the hierarchical whole—are social in their substance. They set up the world in which we live, and their impacts register in the built

environment. Data profiles are the imagined persons for whom the built environment is built. They flourish in the datascape. Indeed, each of us is transformed into profiles thousands of times every day. But they're informational entities, not individuals.

In contrast, our cyberpersonas make us public as highly articulated individuals and do so not through surveillance but through our voluntary self-disclosure. Unlike our passive contributions to the data profile, mostly the electronic traces of our shopping and buying, we actively create our cyberpersonas through the content we publish and through the connections we make with others. The content we publish is our online presence. What others think of our content's quality, expressed by their links and clicks to that content, is our online reputation. Presence and reputation make us visible and knowable as individuals.

This persona-generating self-disclosure has three substantial benefits for the user. Self-expression is one. Anyone can now easily and inexpensively publish content they create, including content whose narrow appeal would not meet marketplace requirements. (Many use the opportunity for personal updates, gossip and sociable chatter, but some, a large number in absolute terms, publish pro bono technical treatises, thoughtful commentaries and other meaningful content on substantive topics.) The second benefit is the change in the relationship of reader and author. Our cyberpersonas are in part co-created with others in the give and take of online conversations. In this process, readers are also authors and vice versa. Third, users can ignore official worldviews. By inserting links of their own choosing and applying tags of their own devising, they can organize and present content, user-generated and otherwise, according to their own worldviews. Our persona-generating activity makes use of all three.

Much of this activity focuses on our lives as consumers. This is explicit by definition in the user-managed interfaces for the "virtual consumer" and at the product review sites,

but it's pervasive across the Web 2.0 terrain. Our social network profiles, personal blog posts and a large portion of the images we upload to file-hosting sites are filled with our tastes in music, books and movies, with details of our hobbies and vacations, with snapshots of the concerts and sports events we go to, with links to the bars and restaurants we like, and so on. So, too, among the collective intelligence applications, most rating, ranking and recommendation systems are designed to filter the vast inventories of cultural commodities available for our consumption.

This emphasis is predictable and unobjectionable. Everyone's everyday life involves a good deal of shopping and buying, and each of us finds some self-expression in what we buy and how we consume what we buy. What's remarkable about the cyberpersona is its highly individuated content, and all the services and applications through which we generate that persona are designed to facilitate it. The virtual consumer interfaces assume we are not merely willing but eager to make our particular wants and needs known to the marketplace. In the Web 2.0 terrain, the social network sites build in functionalities that enable us to share our friends, tastes, interests and activities. Blogs enable us to share our thoughts; the software envisions a personal journal, a window onto the mind of the blogger that's open to others. The images we post to file-hosting sites are literally our views of the world, just as the tags we use at bookmarking sites express our particular worldviews. The collective intelligence applications do not by definition result in self-disclosure, but even their outputs require that we reveal our preferences and predictions to their processing machinery.

Our self-disclosure via these applications—biographical information on our schooling, careers and family lives, updates on our hobbies, indicators of our tastes, preferences and loyalties, snapshots of the public and private events we attend, the topics that interest us and our opinions about them—makes privacy defenders cringe. Their paradigm

assumes we want to withhold ourselves from view. Cyberpersona assumes that we want to advance ourselves into view and do so in the highly articulated form of this and no other individual.

So far, users have asserted only two claims to their online representations. One, we want to control its publicity, that is, which portions of our self-presentation are shared with whom and under what circumstances. Two, we want to take both our content and our connections with us wherever we go on the Web. User claims to the publicity and portability of our cyberpersonas are likely just the beginning. As more of our everyday lives occurs or is replicated online, our presence and reputation online will likely evince additional dimensions and dynamics, the control of which will be in the user's interest.

Whatever dimensions and dynamics emerge, advancing any claim as a matter of "right" is always a complicated effort. It may be more effective to advance claims to our cyberpersonas as "wants" that the marketplace can fill. This could work. The Web is still a fluid medium, very much a work in progress, and Web developers have proven adept at putting users in the driver's seat and letting user behavior guide product evolution. Just as important, many more users will be confronting these matters head-on. As more of our everyday lives occurs or is replicated online, our presence and reputation will grow without our active oversight or cultivation. However, to the extent that a coherent and compelling cyberpersona becomes an increasingly important social asset and an increasingly widespread social expectation, each of us will have to take up with self-conscious purpose the tasks of building, managing and making good use of our online presence and reputation. In other words, whether we advance our claims to our cyberpersonas as rights or as wants, they will likely become responsibilities for most of us. There's only one catch: as a hybrid entity the cyberpersona like the data profile has features with which we are unfamiliar.

Post-Human Entities

Although data profile and cyber-persona make us public in different ways, both are hybrid entities. They are part human. Carbon-based individuals provide the raw material in continuous flows for both. As we change, our simulations change in tandem. They are part machine. The machine determines what we can put into it and how whatever we put in is then served up to others.

The datascape is a mathematical world, and the profile is made of data, comparable and combinable values on common attributes that are processed into probabilities. Cyberspace is a network world, and the persona is made of connections, the links and clicks that we make to others and they make to us. Probability and pattern are the forms in which human raw material appears inside these two machines. Both have features that don't seem to fit the humans from whom they are derived. While the individuated being that each of us experiences is a continuous, whole and bounded entity, our simulacra are contingent, relative and open to others. Indeed, our machine appearances are hard to pin down.

Individuals persist from one moment to the next; that continuousness is our internal subjectivity and our external biography. In contrast, profile and persona do not even exist until they're called up. In the datascape, our attributes sit idle until someone queries the database. In cyberspace, content by or about us sits idle on a server until someone links or clicks to it. To be sure, their life spans differ. The profile is ephemeral; thousands are churned out and thrown away daily. The persona is forever; one's online fragments, wherever they appear, never disappear. Neither profile nor persona can conjure up themselves, however. Someone else must call them up.

Both simulations are always partial, because in every instance they are called up relative to the observer. In the datascape, the query determines the profile. The mind's eye of the beholder envisions a desired end state. The machinery

then sorts through our attributes, and its mathematics call up into the profile only those attributes that contribute in a statistically significant way to the end state envisioned by the query. Thus, every profile portrays us as an object of desire defined by the eye of the beholder. Similarly, cyberpersona is partial because relative but in a network way. Each of us has digital fragments by or about us in various locations across the Web, but the Web is not a lattice of evenly spaced nodes. It's an irregular mesh of knotted and clustered connections, created by the preferential attachments of its users. Which of anyone's fragments become visible to whom depends on the observer's location in the network and the links he encounters and can traverse from that location. Unless the user claims and gathers all his fragments in one spot, our cyberpersonas are not visible in toto. They too are always partial because always relative to the location of the beholder on the network.

Finally, both simulations are open to others. The ability to combine and divide data enables decision makers to expand and narrow the scope of human affairs about which they assert knowledge. In the datascape, they continually reconfigure us with varying others to optimize our utility to their desired end states. The data profile has boundaries, but they are always provisional. Cyberpersona is open in its way; it depends on the "other." The virtual consumer requires the online merchant; they co-create. The social network profile requires friends for its existence; in fact, the activities of those friends provide most of the content of the profile owner's page. So, too, bloggers are visible only to readers; indeed, blogging becomes dynamic only when readers comment, turning posts into conversations. Similarly, whatever collective intelligence is created by those applications that enable us to think together is by definition co-created by their users.

Contingent, relative and open to others, our simulacra seem different from the continuous, whole and bounded individual from whom they are derived. Among the many

ways one could characterize this divergence, three contemporary perspectives are useful to spell out: the post-modern, the post-humanist and the post-human. They take different paths to the same place.

The post-modern view dismisses the continuous, whole and bounded self as a modernist fiction. The internal world each of us knows subjectively has always been a social construct, created from within language and our situations. In this view the hero of modernism including the individuated “in here” has always been contingent, relative and open.

The post-humanist view also sees this autonomous being as a modernist fiction with strengths and weaknesses, cultural and scientific, all of which we can now move beyond. Instead, we should understand the self as including and even give preeminence to our permanently partial identities, the contradictory standpoints from which they appear and the continuous co-authoring of them through interacting with others. These attributes of the more broadly conceived post-humanist self are similar to those proposed for the more narrowly conceived post-human entity with one obvious difference: the post-human view doesn't bother with the hero of modernism.

From a cybernetic perspective, the post-human is an informational entity, continually nourished by the carbon-based world but disembodied and shaped for silicon-based worlds. Transposing human affairs into data is how we are disembodied and move from carbon into silicon. The resulting entity inside the machine lacks the boundaries, coherence and destiny of singular individuals. Instead, it is composed of heterogeneous components that are subject to continuous construction and reconstruction and that manifest different identities under different perspectives.

This cybernetic concept captures at a high level the contexts in which the simulations discussed here arose, the features they have in common and our relationship to them. They are not our descendants, successors or representatives; rather, they are simulations of us, nourished by our lives but

shaped and served up to others by the machines in which they appear. For all the undeniable benefits that data profiles and cyberpersonas deliver, these simulations are human-machine hybrids and as novel entities merit our wary attention. As Norbert Wiener, the founder of cybernetics explained at its outset, “What is used as an element in a machine is in fact an element in a machine.” That's both an observation and a caution.

Notes

1 On the benefit of avoiding obligations, see Thomas Erickson, “The World Wide Web as Social Hypertext” at http://www.pliant.org/personal/Tom_Erickson/SocialHypertext.html. Anonymity also removes the “shadow of the future” from our interactions with others. This concept refers to the following behavior-governing expectation: I act toward you today in the expectation that you will act toward me likewise tomorrow

2 Joe Arena, “Framing an Ideology of Information: Retail Credit and the Mass Media, 1910-1930,” *Media, Culture and Society* vol. 18 (1996), pp. 423-445 explains retail credit extracted the individual's trustworthiness from local social networks and quantified into a number, making it portable and comparable.