LA RB

The Digital Revolution: Debating the Promise and Perils of the Internet, Automation, and Algorithmic Lives in the Last Years of the Obama Administration

LOS ANGELES REVIEW OF BOOKS | DIGITAL EDITIONS

EDITORS: MICHELE PRIDMORE-BROWN AND JULIEN CROCKETT

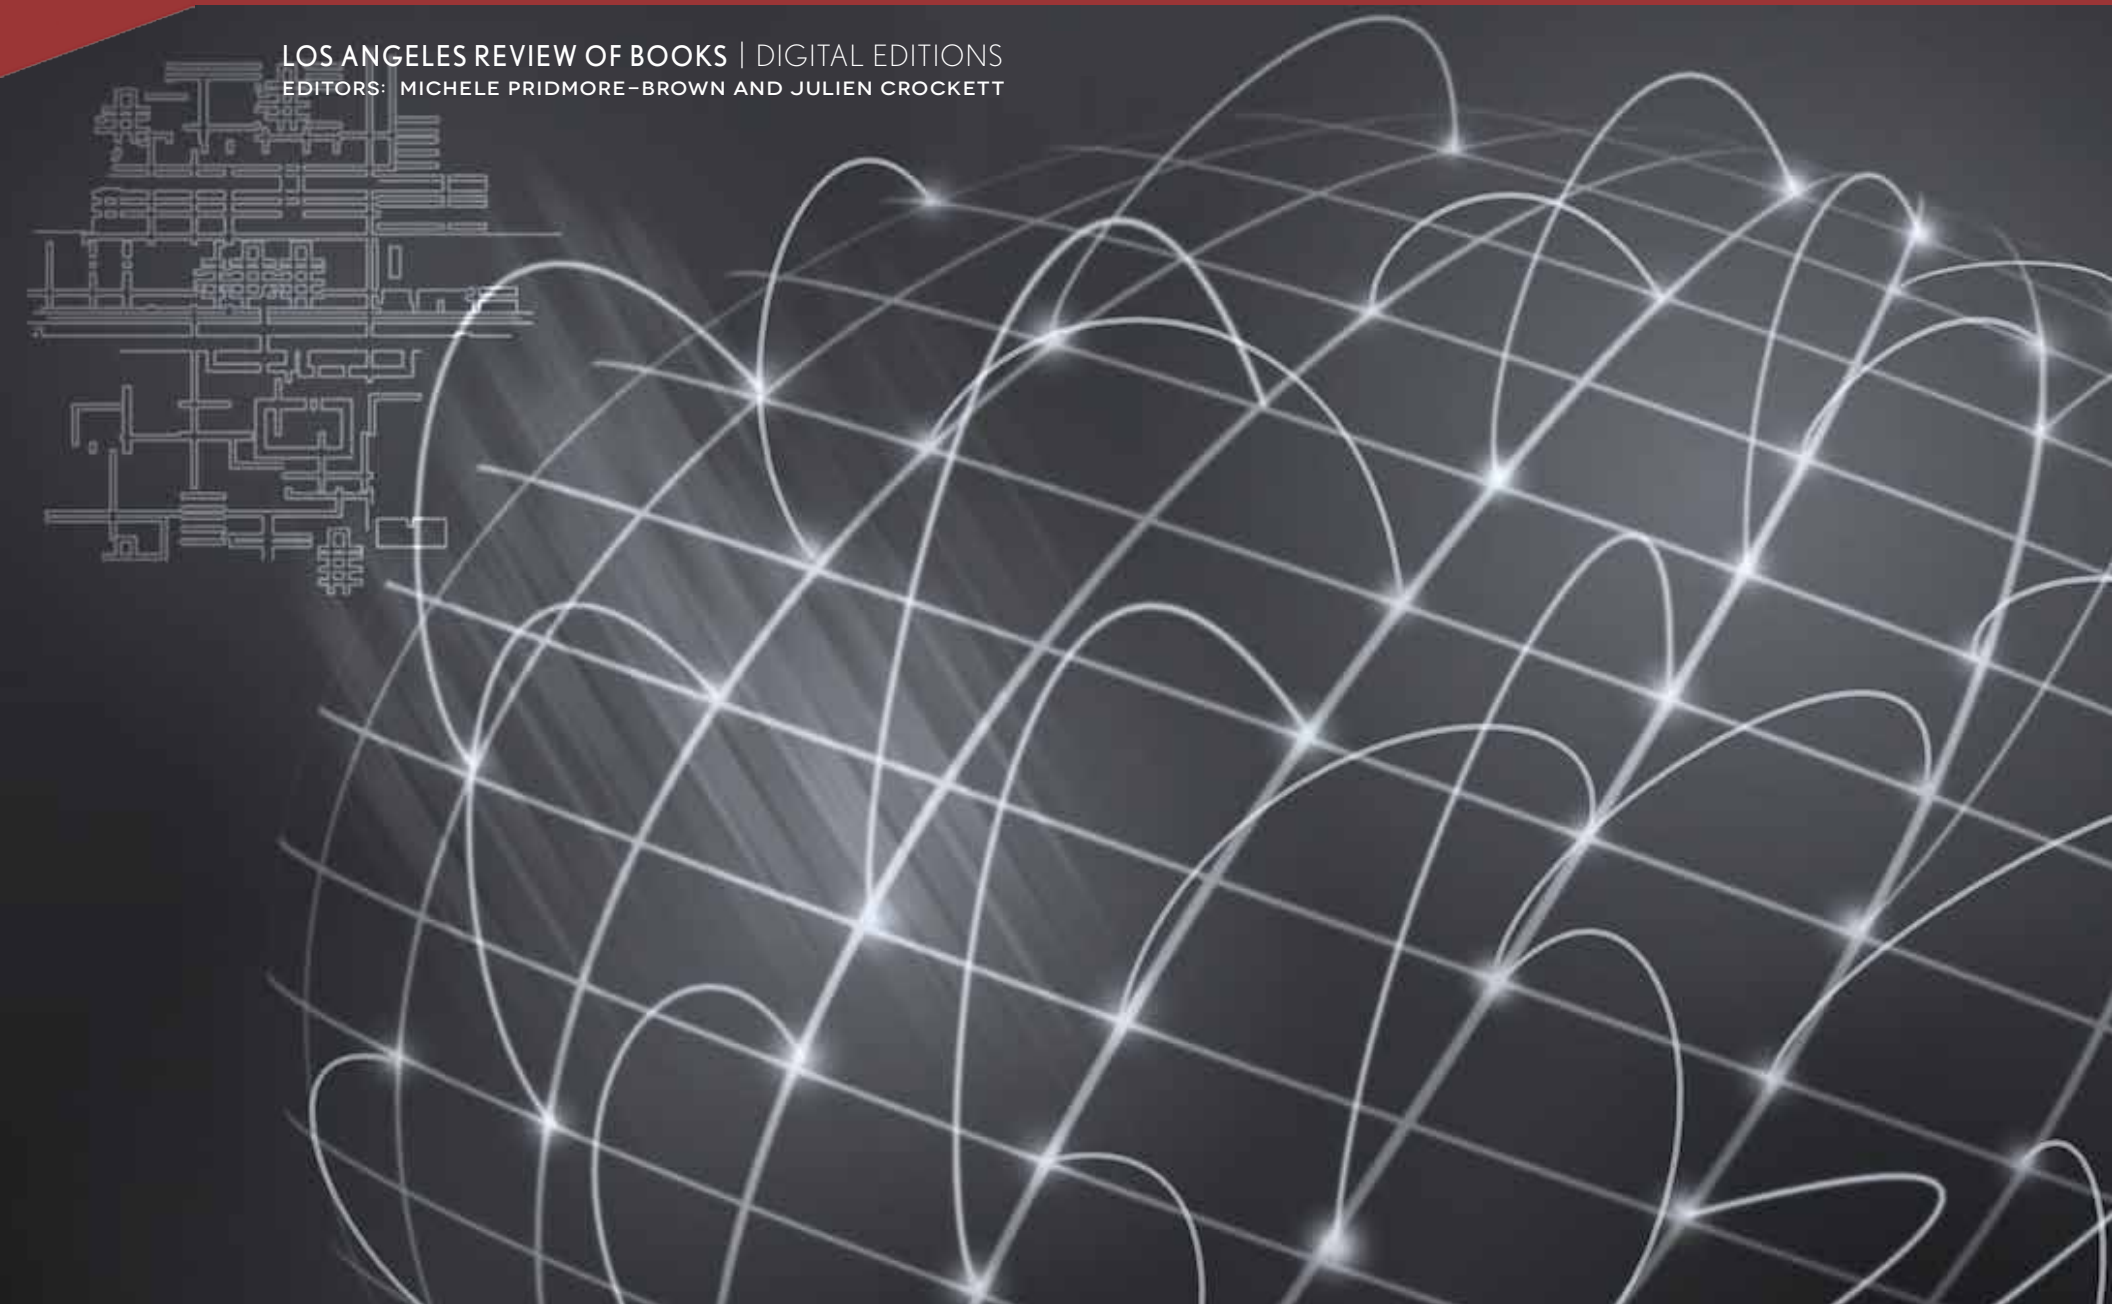

cover image courtesy of solidware.io/

Keywords: Algorithms, Artificial Intelligence, Automation, Big Data, Contextual Integrity, Cyberculture, Cyberspace, Digital Revolution, Dataism, Data Protection, Disruption, Echo Chamber, Efficiency, Emblematic Artifacts, Expository Society, Gamification, Governance, Internet, Internet of Things, Internet Privacy, Machine Learning, Networked-knowledge, Obfuscation, Regulation, Robotization, Social Contagion, Silicon Valley, Social Engineering, Social Network, Sociotechnical Ecology, Surveillance, Techno-Optimistic Bubble, Techno-Utopian Bubble, Turing Machine, Uncanny Valley.

TABLE OF CONTENTS

Introduction

Plenty of rapturous claims have been made about the internet as an agent of democratization and innovation. Many more claims have been made about its Pandora-like perils, and how these herald our individual and collective downfalls, whether by taking away our jobs, exposing our darkest selves, or turning us into mindless automata disappearing into "click-bait rabbit holes." The internet has allegedly enabled a new kind of populism, as well as political gridlock and the infelicities of the 2016 US presidential election. Clearly, it is the stuff of paradox — and possibly, as historians like Yuval Harari and others suggest, of Faustian bargains.

Like the commercial and industrial revolutions before it, this particular revolution opens up social spaces that alter how we affiliate and think. On the plus side, we can now communicate and connect with others like never before — not just with the approximately 150 embodied others whom anthropologists contend most people knew in the past, but with thousands and indeed millions of digital others. But it's much harder to trust them, and to trust "information." It is easier to "hide" — and also harder, which, depending on where you sit, or lurk, is a good thing, or not. Surveillance is more total, even while it's more self-inflicted — and pleasurable; we love Netflix "knowing" us so well, and so we happily expose more of ourselves to get more of what we want. This is part of the Faustian bargain.

The LARB Science and Technology section has been capturing scholarly and popular views on the digital revolution in a series of essays and reviews that insist on historical perspective — on the longue durée. They were all written in the last years of the Obama administration, before the Trump one; they were first published on the LARB website and are now collected in this volume. They express what experts in their respective fields got right — and what they may have gotten wrong. They examine the stakes. In some cases, our contributors dismantle their colleagues' arguments, especially when those arguments express a certain knee-jerk zeitgeist (e.g., the digital age is deskilling us or making us stupid). Internet philosopher David Weinberger of Harvard University, for instance, takes on the argument that the net is turning us into passive knowers. On the contrary, he counters, the net is transforming knowledge in ways that reveal the flaws inherent in past ways of knowing. "Networked-knowing" is, in his view, a positive phenomenon — it replaces the manufactured or "curated cohesion" of past knowledge regimes. As for the claim that the net reinforces echo chambers (and false news), he plays the contrarian again, countering that those chambers are now, thanks to the net, shot through with holes that anyone, including a teen-ager trapped in an otherwise airless cult, can follow just by clicking her finger.

The fact that no echo chamber is impervious to "information" from the outside certainly seems like a good thing. The fact that there's quite possibly no algorithm or security system that isn't potentially "leaky" or hack-able is perhaps a less good thing, but part of the same digital coin. Our purported doom and salvation lie in the same places.

In an essay he wrote for LARB in 2014 and included here (see "The Manipulators"), journalist and blogger Nicholas Carr, author of several influential books on technology and culture, claimed that we may look back on 2014 as the year the internet "began to grow up" and the public was called upon to guide the technology rather than the other way around. He made this claim seven years after the launch of the iPhone and "tech boom" of 2007, as identified by journalist Thomas Friedman — when Facebook, Twitter, and "the cloud" took off, and when, in Friedman's opinion, connectivity and computing got so fast and cheap and ubiquitous that they vastly outpaced social institutions. Writing in 2014, Carr thought that "we" would concertedly rebel against so-called corporate surveillance creep, and against being lab rats — as in the OK Cupid and Facebook social engineering experiments of that year, which he addresses in his article here.

Clearly, we haven't rebelled. Or, to use Carr's metaphor, the internet has not "grown up" — and neither have we, its creators and users. On the contrary, "it" has only gotten better and more efficient at pressing our buttons.

In his 2016 book *The Master Algorithm*, computer scientist Pedro Domingos emphasizes that, in the digital age, "growing up" means knowing our algorithms. A responsible citizen is, in other words, a digitally savvy citizen. But, as sociologist Michael S. Evans counters in his review entitled "Algorithms: The Future that Already Happened," we can already be managed — tracked, nudged, fixed — faster than we can respond. In addition, being technically knowledgeable simply isn't expedient in our daily lives. Nor is "obfuscating" our digital tracks (in this regard, see Evan Selinger's review: "Internet Privacy: Stepping Up Our Self-Defense Game"). What we really need, quips Michael Evans, is a "data reset button" — so that we can start over every few years on a digitally fresh playing field in the manner of a 5-year plan or Biblical Jubilee.

A pause button might be more feasible if equally unlikely. But Evans has a point even if he is half-joking. And, of course, that's another paradox of sorts. In the digital world, there is no absolute forgetting — at least for now, even as, on an individual level, we are forgetting more and more in a process dubbed "cognitive offloading;" the internet is of course not just our aide-mémoire, but transforming our memories and how we think. Is it transforming us even more profoundly than the commercial and industrial revolutions did?

An impossible question perhaps, but in his illuminating essay entitled "Algorithmic Life," Berkeley historian and sociologist of science Massimo Mazzotti suggests, like Carr above, that we're at a threshold moment. As happened with the word "technology" in the middle of the 20th century, our uses of the word "algorithm" are capturing new and unexpected processes of change. He explains how algorithms naturalize ways of doing things. They naturalize specific cultures. They're tools for doing, as he puts it, and tools for thinking. For these reasons he calls them "emblematic artifacts" that, like the clock and computer before them, shape how we understand ourselves and act within the world. They're like "powerful rhetorical arguments" in their ability to create and normalize social worlds. For Mazzotti, being a grown-up means understanding their "sociotechnical ecologies" — even as, thanks to innovations in machine learning, they increasingly defy our ability to comprehend them. It means asking questions about "the process," and understanding whose worlds they're making and whose biases they're reinforcing.

In other essays and reviews, contributors take on particular aspects of the digital revolution. Literary scholar Dennis Tenen of Columbia University addresses the voluntary surveillance paradox alluded to earlier with respect to Netflix — but he uses the prison context as a point of departure. Legal scholar Frank Pasquale, author of the important 2015 book *The Black Box Society: The Secret Algorithms that Control Money and Information*, denounces the "false promise" of "disruption."

In a new afterword to *TechGnosis*, published with LARB in 2015 under the book's subtitle "Myth, Magic and Mysticism in the Age of Information," the genre-defying writer Erik Davis reflects on "the contemporary urge to 'gamify' our social and technological interactions," and deconstructs net-enabled "enchantment" cum "weirdness." Berkeley information scholar Geoff Nunberg addresses Nicholas Carr's latest book *Utopia is Creepy* (a compilation of his essays and 2005-2015 blog posts) to plumb the technological uncanny, but insists it's in fact banal. He describes how, in 1964, he took a year off from college and worked in the General Motors Pavilion at the New York World's Fair escorting VIP guests on the Futurama II ride. Humans were stick figures in the new technological corporate-sponsored imaginary. For Nunberg, we're those stick figures — even as we're ostensibly being ever more empowered and enhanced. "What's most striking," about the sensor-saturated world of the future, he observes, isn't just the creepiness of new devices that watch and sense our every move, but "how trivial and pedestrian they can be." The future isn't more exciting, just more

"efficient," he concludes.

Professor of Culture and Media at The New School McKenzie Wark addresses the artist's escalating practical dilemmas in an age when new digitally-enabled entities are taking an ever growing cut of the creative pie — at the very moment, in other words, when unsponsored against-the-grain creativity seems most important.

In separate essays, both the aforementioned legal scholar Frank Pasquale and philosopher Evan Selinger address automation's likely effect on jobs; Pasquale expertly dismantles some of the hype around a jobless future and Selinger thoughtfully considers, by way of a review of Martin Ford's 2015 book *Rise of the Robots: Technology and the Threat of a Jobless Future*, whether a techno-optimistic "bubble" is indeed, as claimed by Ford, impeding smart discussion of the perils of automation.

Contributor Eli Mackinnon uses the occasion of the republication of Andrew Hodges's biography of Alan Turing to look back on the "jailbreaking" life that spawned the digital revolution. Another contributor, Julie Carpenter, asks what's "fair" in war. A human-robot interaction researcher, she looks to both past and future, examining how age-old moral dilemmas are amplified when fleshly bodies meet digital realities in 21st century militarized spaces.

Collectively, these essays thus address the digital revolution from a variety of disciplinary and ideological perspectives. They extol and debunk. They query, rather than take for granted, terms like "disruption," "echo chambers," "algorithm," "efficiency," "gamification," "threshold moment," and even old standbys like the words "revolution" and "technology" — all of them the currency of the digital age. They make legible the entanglement of societal values with technologies, as well as the eternal return of the same motifs (progress, utopia and dystopia, emancipation, millenarianism, or the term Faustian bargain referenced earlier). They look to historical antecedents, in other words, for what seem like de novo developments. Perhaps most importantly, they also, at least in some cases, dare to think about how we might, while we still can, go about shaping our digital futures. To close with Mazzotti's point at the end of his essay: algorithms are now "the doors onto our futures;" we should at least be self-conscious about which ones we're opening.

Michele Pridmore-Brown
Julien Crockett

January 20, 2017.

1

Rethinking Knowledge in the Internet Age

by David Weinberger (May 2, 2016)

David Weinberger counters philosopher Michael P. Lynch's claim in The Internet of Us: Knowing More and Understanding Less in the Age of Big Data *that the internet is impeding our ability to learn and know in complex ways.*

THE INTERNET started out as the Information Highway, the Great Emancipator of knowledge, and as an assured tool for generating a well-informed citizenry. But, over the past 15 years, that optimism has given way to cynicism and fear — we have taught our children that the net is a swamp of lies spun by idiots and true believers, and, worse still, polluted by commercial entities whose sole aim is to have us click to the next ad-riddled page.

Perhaps our attitude to the net has changed because we now see how bad it is for knowledge. Or perhaps the net has so utterly transformed knowledge that we don't recognize knowledge when we see it.

For philosopher Michael P. Lynch, our fears are warranted — the internet is a wrong turn in the history of knowledge. "Information technology," Professor Lynch argues in his new book, *The Internet of Us*, "while expanding our ability to know in one way, is actually impeding our ability to know in other, more complex ways." He pursues his argument with commendable seriousness, clarity, and attunement to historical context — and yet he misses where knowledge actually lives on the net, focusing instead on just one aspect of the phenomenon of knowledge. He is far from alone in this.

The net in fact exposes problems that have long lurked in our epistemology, problems that come into stark relief when knowledge is freed of paper, and we freely connect with it and through it across all boundaries of time and place. There's something about how we've been thinking about knowledge — something inherent in traditional epistemology — that blinds Lynch and many others to the knowledge-enhancing aspects of what's happening on the screens in front of us.

Knowing Beyond Google

The rhetoric of internet criticism often follows a typical pattern, such as, "Yes, but," as in: "The internet does many wonderful things, but …," followed by a long list of everything that's wrong with the net. Lynch is far more thoughtful than most, but his arguments nonetheless suffer from the flaws that typically bedevil this type of criticism. His idea of what the net gets right — the "Yes" before the "but" — is itself qualified, and a

good place to start. He characterizes knowing on the net as "Google-knowing," or what we might otherwise call "looking something up." He acknowledges that having instantaneous access to facts via Google-knowing ranges from handy to lifesaving, but his conception misses the full picture in two important ways.

First, we should include in Google-knowing more than just factual look-ups. If, in the 19th century, almanacs commoditized facts, then the net has taken the process one big step further, commoditizing the encyclopedia article. Not only can we look up when Thomas Jefferson was the United States minister to France (May 1785-September 1789), we can also get some quick context about what that meant to the Colonies, to France, and even to him. (To his credit, Lynch does not engage in the usual Wikipedia bashing.)

More important, to focus a discussion of internet knowledge on people looking things up on Google, or even in Wikipedia, is like describing libraries as places where people use the encyclopedia. Far more typically, when we're on the net, someone or some service links us to a bit of news that's interesting to us. It might be politics, quantum physics, or the Kardashians. Whatever it is, let's say we want to know more. Before the era of the net, the reader's curiosity was bounded by the physical rectangle in the newspaper within which the article sat. On the net, if the source of the tidbit — a tweet, a Facebook post, an email — doesn't itself contain links to additional information, then we can pop some terms into a search engine and find more avenues to explore; push a comment or question back into the social medium through which we first learned of the topic; hop it over to a different social network; or, even reach out to the thinker or writer who stimulated the discussion in the first place.

This is the phenomenon of knowing these days. Surprisingly, it does not figure much in Lynch's book except in examples of how knowledge goes wrong. "Sometimes we need to know more than the facts; sometimes we want to understand," he asserts, before concluding that "it is our capacity to understand that our digital form of life often undersells, and which more data alone can't give us." That's a trenchant criticism of the net, but only if we assume that on the net, all knowing is Google-knowing.

In fact, knowledge is now networked: made up of loose-edged groups of people who discuss and spread ideas, creating a web of links among different viewpoints. That's how scholars in virtually every discipline do their work — from their initial research, to the conversations that forge research into ideas, to carrying ideas into public discourse. Scholar or not, whatever topic initially piques our interest, the net encourages us to learn more. Perhaps we follow links, or are involved in multiyear conversations on stable mailing lists, or throw ideas out onto Twitter, or post first drafts at arXiv.org, or set up Facebook pages, or pose and answer questions at Quora or Stack Overflow, or do "post-publication peer review" at PubPeer.com. There has never been a better time to be curious, and that's not only because there are so many facts available — it's because there are so many people with whom we can interact.

While *The Internet of Us* mentions many of these new formations, it is usually to dismiss them: "Social media and the blogosphere are filled with 'reasoning'" that's either "blatant marketing or aimed only at supporting what people already believe," writes Lynch. He deals with MOOCs, or massive open online courses, as if they've already found their final form and as if MOOCs are all the net has done for education. There is only a single mention of the Open Access movement late in the book. Although Lynch concedes that data analysis in the digital humanities can sometimes yield interesting results, that admission comes only after he has spent several pages demonstrating in detail the flaws in the sole digital humanities article he uses as his example.

Why does Lynch miss so much? He has written an intelligent book that struggles honestly with important questions: Is the net turning us into passive knowers? Is it degrading our ability to reason? What can we do about this?

As our friends on Reddit might put it: He misses it because classic epistemology.

Setting the Rules of the Knowledge Game

Lynch acknowledges from the start that he's writing as a philosophy professor embedded in the classical tradition: "My aim is to examine the philosophical foundations of what I'll call our digital form of life," he declares, using an apt Wittgensteinian phrase.

He thus does not claim to be breaking new ground in epistemology so much as applying the wisdom of "The Tradition" to our digital world. He writes for an undergrad rather than specialized level, and introduces concepts as they become relevant. For instance, he nicely explains the Socratic distinction between right opinion and knowledge: knowledge entails justifying one's opinions. This throws him straight into the question that has long occupied traditional epistemology: what constitutes justification?

His answer is also traditional. He starts with the lowest level of knowledge: the type of knowing that we share with other animals. This "receptive knowledge" springs straight from sensation: "if our visual system, for example, is working as it should, and our representation of the world is accurate — if we see things as they are — then we come to know" in this basic way. True, like so many others, he modernizes the vocabulary, using phrases like "downloading" information, but his description nonetheless constitutes straightforward empiricism.

He then climbs the empiricist's ladder. How about knowledge that is more abstract or general — the knowledge that can, for instance, be found on the net? He dismisses the idea of the heroic single knower, acknowledging that "reflective" knowers — a term he pairs with "responsible" — need to rely on others. This reliance requires learning and applying the rules of reason in public spaces, where, heeding Kant, we must think for ourselves and treat others as equally autonomous agents. But, suppose at bottom, all chains of reasons end with mere belief?

In a rigorous and frank chapter, Lynch faces this question, which has dogged epistemology from its origin, and which he believes the internet is making "even more pressing." Invoking Wittgenstein again, he writes that, in order to be reasonable, we have "to be willing to play the game of giving and asking for reasons …" This edges him dangerously close to the postmodern world in which such games cannot confidently assert their privilege over other games. Or, if you call this a "paradigm" instead of a game, then you are taken to the territory of Thomas S. Kuhn's 1962 *The Structure of Scientific Revolutions*. Or to a Heideggerian critique of rational scientific knowing as one way of disclosing the world, which needs to be grounded in a more basic analysis of disclosure itself. Or, perhaps, to a Foucauldian examination of knowledge as a historical expression of power. Or to Bruno Latour's explication of the dependence of knowledge on its instruments and the networks in which they reside. Or to the "embodied mind" theory put forward by Francisco J. Varela, Eleanor Rosch, and others that takes the mind to include the body acting in the world.

None of these are necessary paths to take. Some may be silly. But Lynch takes none of them, declaring that "the game of giving and asking for reasons" is to be played "by rules most of us could accept were we to stop to think about it." He stays on the Enlightenment side of the divide, well clear of the postmodern epistemic world.

Indeed, that thoughtful people would come up with the same rules remains the great hope of the democracies spawned by the Enlightenment. It buttressed Barack Obama's patient insistence that opposing sides can find common ground. It is at the heart of Al Gore's *The Assault on Reason*, an assault he believes the internet can help counter, although it was published in 2007, before the current internet backlash had taken over the discourse.

But if the internet has taught nothing else, it has taught us that we will never agree about anything. For every fact on the internet, there is an equal and opposite fact. It doesn't help that not every purported fact is in fact a fact, or that facts tend to have multiple layers of truth. Facts simply are not going to play the role in building consensus that we had hoped.

This would be a source of pure despair if the internet were not also enabling us to see that before it existed we never agreed about anything either. Before the net, what we read and saw was so tightly controlled by cartels of well-intentioned professionals that dissenting voices were barely heard. True, many of those dissenting voices were wrong and sometimes they were spouting lunacy, but we marginalized all but the one percent of the epistemically privileged. We achieved great science but at a high price throughout the rest of the cultural landscape, and sometimes within science, too.

This fragmentation of knowledge is a fact that knowledge cannot overcome. How, then, do we best live with it? How do we flourish now that we can't reason ourselves back together?

The Sound of One Echo Chamber Clapping

First, there is little cause for panic or despair. Knowledge has always been a local product. Now, with our localities connected globally, we are more clearly seeing what has always been the truth of our condition.

And it's not as if we are oblivious to this issue. Handling the challenges posed by networked knowing is at the core of almost every website and internet service. Each site or service takes steps to establish the rules of engagement so that users can come to appropriate levels of trust about the claims made by the site and by other users. For Uber, it's a ratings and comments system for drivers and passengers. For Public Library of Science (PLOS), it's a peer review process and post-publication comments. For Reddit, it's the up- and down-voting of vigorous and extended discussion threads. None is perfect, but then knowledge has never been perfect. Each is, however, appropriate for its domain.
But, instead of flawed, domain-specific ways of knowing on the net, Lynch sees in-grown, hermetic clusters of people who simply confirm each other's baseless beliefs and drive themselves to extremes. In short: He sees echo chambers, which Cass Sunstein notably discussed in his 2001 book, *Republic.com*.

Echo chambers are indeed a problem. Large commercial sites abet them, making money by getting us to click on links to yet more ad-laced pages. And, of course, these sites exploit the fact that we're more likely to click on a link to an idea with which we agree than one we hate. Even respectable news organizations have learned that clickbait headlines work best if they play into our existing beliefs. "Filter bubbles," as Eli Pariser calls them, pay.

But thinking of the net as a set of echo chambers is just bad phenomenology. If you're using email, Twitter, or Facebook, or if you do your Google-knowing via a search engine, you are inevitably face to face with at least some ideas and links that don't come from your particular echo chamber. Plato's Cave is now shot through with holes that transmit light, sounds, and smells that you can follow just by clicking your finger.

Do we have to follow them? Of course not. We tend to follow those that make sense to us. Our interests turn out to be distressingly parochial and local in most cases, as Ethan Zuckerman demonstrates in *Digital Cosmopolitans*. But, since epistemologists like Lynch insist echo chambers have degraded knowledge, we have to ask what our media diet consisted of before the net.

The answer isn't pretty: Three channels of nightly news; a small handful of newspapers in major cities; a nice spread of national magazines, each one its own echo chamber; a Great Books series launched in 1952 that

consisted of works by 130 authors, not one of whom was a woman or black, and almost all of whom were within the European tradition.

Of course, the decades after the 1950s offered a greater diversity of voices, but the basic mechanisms of knowledge continued to do the work of echo chambers: identify and publish the work of experts, vet knowledge to exclude what does not fit, and create a coherent body of knowledge that becomes harder and harder to challenge precisely because of its curated cohesion. Certainly the content was incomparably better than what comes out of toxic internet echo chambers. But the problem with echo chambers has to do with their structure, and traditional knowledge comes from that same structure.

The net is making clear how important "echo chambers" are to knowledge and even more so to understanding. If you care about molecular gastronomy and hear about a new technique, you'll go to your favorite molecular gastronomy sites to learn more. If you're a supporter of Net Neutrality and there's a court ruling you don't understand, you'll go to a site that shares your values to get the explanation. If you are a feminist and a new pay equity law passes, you're not going to go to a male supremacy site to find out what it means for you. Knowledge and culture depend on like-minded individuals joining together and iterating over tiny differences. This is how the net works. This is also how traditional knowing works. We did not like to acknowledge that. Now we can't avoid it.

The Medium is the Knowledge

Perhaps our chief epistemic avoidance mechanism was turning knowing into the production of a type of content — knowledge — that we convinced ourselves had to be independent of the knower in two senses.

First, we devised methodologies that try to keep the vagaries of the individual out of the process of creating knowledge. The scientific method works. Journalistic objectivity continues to be reevaluated. (I should note that on the question of objectivity Professor Lynch contrasts his views explicitly with mine, as he does when he considers networked knowledge. He does so without animus, and I certainly respect and appreciate his criticism.)

Second, we physically separated knowledge from individuals by externalizing it (e.g., books). What started in Greece as a particular class of belief became a body of printed statements that could be called knowledge even if there was no one left to believe them. Obviously, this has been wildly successful for our species, but it also meant that the medium of externalization — paper — has shaped knowledge to fit its peculiarities. Knowledge is settled belief because ink settles into paper. Knowledge divides into topics because books can only be so big. Our most revered form of knowledge consists of long chains of reasoning because books are sequential and not suitable for optional digressions. Knowledge is something most of us read passively because books are inert — just as Socrates had opined in the *Phaedrus*. Knowledge is a type of content independent of believers because books are physical objects that outlast their authors and readers.

There is nothing natural about knowledge, so it is not surprising that it shaped itself to its medium of storage and communication. But now we have a new medium, and knowledge is taking on its properties. That is why looking for traditional knowledge on the net leads us to miss the phenomenon of knowledge there. It's also why the idea of knowledge is not much discussed on the net. It's possible that, in the next generation, the noun "knowledge" will become old-fashioned and, in the generation after that, archaic.

That wouldn't mean that we have given up on the project of knowing our world. Instead, it would mean that that project has changed shape — from content to networks that mix ideas and sociality. That's already happening. Knowledge is becoming what happens when links connect differences and people.

This is new but it should also sound quite familiar, for it is very close to how scholarship works in the traditional world: people who care about a topic argue about it among themselves. Networked knowledge opens up the discussion and the participants remain linked and engaged, not always expecting final resolution. Knowledge exists in the linkages of difference and disagreement.

The net is demonstrating the weakness of knowledge as finished, settled, and static content. It's doing so by plunging us deeper into knowing.

The Networking of Knowledge

Lynch disagrees. He sees more misinformation than knowledge on the net, and he believes most of us are abdicating our responsibility for sorting the one from the other. Google-knowing is, for him, akin to consulting an unreliable almanac. That may be an inevitable consequence of scaling up almanacs to include answers to just about every question imaginable, although I suspect that Google does quite well with questions concerning information literally in almanacs. Curated sites such as WolframAlpha.com undoubtedly do even better.

But let's grant Lynch his concerns. They are real even if they are not the whole story. So, what do we do about them?

Once again Lynch's answer relies on a traditional assumption that our experience of the net calls into question: agreeing on our "epistemic principles" because if we're not playing by the same rules, then it's "game over," as he puts it. But which rules? Lynch tells us that no matter what we believe in our private lives — "the Bible, or the Koran, or *Dianetics*" — we should "support institutions like the National Science Foundation or the National Endowment for the Humanities …" for they are "*doing the work of democracy*." (Emphasis in the original.)

Like everyone else I know who packs their Whole Foods groceries into NPR tote bags, I, too, like those institutions. But if I am a committed Christian, Muslim, or Scientologist, those religious beliefs are at the heart of what I *know*. Telling me to put those beliefs aside so I can engage in a reasonable discussion is to ask me to play by someone else's rules. And the problem is repeated if we accept that it's not just religious people who disagree about the rules of engagement for developing knowledge. Arguments over what constitutes acceptable evidence are common in just about every field. (Lynch's condescension about religion might be overcome if he were able to engage with the faithful on their own terms.)

I do agree that when it comes to making public policy, we want evidence-based discussions whose rules do not require private access to sources of truth. Likewise, I want the scientific method to settle questions of science and the jurisprudence system to settle questions of criminal culpability. But the net has proven what philosophers, historians, neurologists, and sociologists have already told us: it's not that simple. Science is deeply embedded in systems of instruments, power, economics, gender, psychology — in short, in all the different facets of being human.

Although we agree that it's good to "promote institutions that encourage cooperation, and even face-to-face contact with people who have very different views," perhaps the way forward seems clearer to Lynch than it does to me for the following reason: beneath his traditional epistemology is a traditional ontology that takes the world as that real stuff that is outside our heads. Citing John Locke and the internet philosopher, Luciano Floridi, he argues that digital objects are more abstract than real things because they are purely formal — they are copies without originals — and thus, he implies, they are less real. He warns us that we shouldn't let the digital world "blind us to the world that is not constructed, to the world that predates our digital selves," the

world of things that are *found* and not *made*, as he puts it. But where is the found thing that we don't make by *taking* it in one way or another? In taking the clock as a timepiece and the bullet as a way to create unasked-for holes at a distance, we are relying on a contextualizing world that has been constructed even before we got here.

Lynch is able to treat the net as a simulacrum because he thinks knowing is something we do in our heads. We build up to meaning by starting with sensation. But the net, in his view, is sensation without a real referent. It is a representation of a representation. It is therefore too bad that he dismisses Andy Clark and David Chalmer's "extended mind" idea by saying, "it might be right but we don't have to go that far" because the mind is already extended, by which Lynch means that we rely on the testimony of others to justify our beliefs. But the extended mind concept says something more: we think with tools. The physicist cannot think about a problem without using a white board. An accountant needs a calculator. The philosopher needs books and writing materials and perhaps a fire and a glass of sherry. We think with tools. We think out in the world, not in inner representations of the world. And now we have new tools for thought. These tools include not just search engines, but everything from web pages to complex multi-modal networks of experts and amateurs. That is where thinking and knowing is now happening.

More to Lynch's point, it is where *understanding* is happening. Lynch acknowledges that we are "cognitively interconnected" by the internet, but, he writes, it's our Google-knowing that's becoming networked, and, as he notes, understanding needs more than that. When he explores the idea of networked knowledge, he uses as his examples the wisdom of crowds as expounded by James Surowiecki, as well as prediction markets. These are important and interesting network phenomena, but they are online crowds in which knowledge is generated by averaging the contributions of individuals. This leads Lynch to his summary dismissal of networked knowledge: "You can't take the individual out of the equation."

No, you can't, and the most important new form of networked knowledge does not. A network in our new age is a set of individuals who maintain their individuality while in relation to one another — that is, while maintaining their differences. The knowledge exists in their individual expressions and, most importantly, in their links to one another. This is knowledge that is not an averaging of inputs, but a maintenance of differences-in-connection.

I suspect that Lynch misses all those places where understanding is happening on the net because he is looking for traditional knowledge that arrives at a definite conclusion, is written down, and serves as a reliable authority. But networks are conducive to a type of knowing that paper makes difficult. Networked, collaborative understanding happens in posts, blogs, messages, tweets, and discussion threads. It even happens in exchanges that occur over the course of 20 years in a closed forum like a mailing list — an echo chamber if you want, or a tiny Republic of Letters if you prefer. It just doesn't have the earmarks of traditional paper-based knowledge.

Lynch sees these new locales of knowledge, but thinks networked knowledge is never resolved because it can't agree on what grounds a chain of reasoning. Lynch counters with the primacy of "the objective world itself," which we can know reliably through the senses. So we are back to empiricism. This leads Lynch to advise us to literally get "up off the couch" and plunge "into the whirlpool of actual experience." "To escape your circle of justification, do what you do with any circle: step outside its borders and breathe in the environment on the outside."

But when you do so, your every perception will further confirm your existing assumptions about how the world works — because understanding is, fundamentally, the act of assimilating the new to an existing context. The inside of your skull is an echo chamber.

There's tremendous value in consulting existing bodies of well-vetted beliefs, and, to their credit, teachers like Professor Lynch expose us to that value. But there is also value in the networking of knowledge in which ideas are linked in their differences. We can go wrong in those networks, but we can also go very right, achieving a new sense of how knowledge goes together even if it never fully coheres.

The networking of knowledge does not achieve the aims traditional knowing has set for itself. It is settled only within a community of believers — and not all communities of believers are right. It is inextricable from its social context. It inevitably contains differences, but those differences are now linked. It is as discursive as the net itself. It often comes in small bites, but those bites are linked out to a world larger than all the libraries that ever existed. Everyone gets to speak, even stupid, evil people. Authority generally has to be earned, not declared. The rules of reasoning vary not only within domains but within each conversational forum. Knowledge is being replaced by knowing, and knowing is now a free-for-all. At its best, this knowing does what Lynch recommends: it thinks explicitly about its rules of justification. At its worst, it's a howling mob.

There is endless evidence to support pessimistic or optimistic views, for both are true. This is the greatest time to be a curious person who wants to learn, and it is the greatest time to be a complete idiot. The net is revealing both the power of our traditional ways of knowing and the fact that traditional knowing has always been a product of flawed humans going wrong and going right together. Knowledge cannot liberate itself from this. Ultimately, knowledge's only hope is for more and better humanity.

¤

David Weinberger, a senior researcher at Harvard's Berkman Klein Center for Internet & Society, writes about how the internet is changing our ideas. He is the author of Too Big to Know: Rethinking Knowledge Now that the Facts Aren't the Facts, Experts are Everywhere, and the Smartest Person in the Room is the Room *(2012).*

2

The Manipulators: Facebook's Social Engineering Project

by Nicholas Carr (September 14, 2014)

Using a controversial Facebook study and a ruling by the European Union's highest court ("the right to be forgotten"), Nicholas Carr argues that the internet has reached a new stage in its development — one in which the public "will be called upon" to guide the technology rather than the other way around.

SINCE THE LAUNCH of Netscape and Yahoo! 20 years ago, the development of the internet has been a story of new companies and new products, a story shaped largely by the interests of entrepreneurs and venture capitalists. The plot has been linear; the pace, relentless. In 1995 came Amazon and Craigslist; in 1997, Google and Netflix; in 1999, Napster and Blogger; in 2001, iTunes; in 2003, MySpace; in 2004, Facebook; in 2005, YouTube; in 2006, Twitter; in 2007, the iPhone and the Kindle; in 2008, Airbnb; in 2010, Instagram; in 2011, Snapchat; in 2012, Coursera; in 2013, Google Glass. It has been a carnival ride, and we, the public, have been the giddy passengers.

This year something changed. The big news about the net came not in the form of buzzy startups or cool gadgets, but in the shape of two dry, arcane documents. One was a scientific paper describing an experiment in which researchers attempted to alter the moods of Facebook users by secretly manipulating the messages they saw. The other was a ruling by the European Union's highest court granting citizens the right to have outdated or inaccurate information about them erased from Google and other search engines. Both documents provoked consternation, anger, and argument. Both raised important, complicated issues without resolving them. Arriving in the wake of revelations about the NSA's online spying operation, both seemed to herald, in very different ways, a new stage in the net's history — one in which the public will be called upon to guide the technology, rather than the other way around. We may look back on 2014 as the year the internet began to grow up.

Spreading "Emotional Contagion"

The Facebook study seemed fated to stir up controversy. Its title reads like a bulletin from a dystopian future: *Experimental Evidence of Massive-Scale Emotional Contagion through Social Networks*. But when, on June 2, the paper first appeared on the website of the *Proceedings of the National Academy of Sciences of the United States of America* (PNAS), it drew little notice or comment. It sank quietly into the swamp of academic publishing. That changed abruptly three weeks later, on June 26, when technology reporter Aviva Rutkin posted a brief account of the study on the website of *New Scientist* magazine. She noted that the research had

been run by a Facebook employee, a social psychologist named Adam Kramer who worked in the firm's large Data Science unit, and that it had involved more than half a million members of the social network. Smelling a scandal, other journalists rushed to the PNAS site to give the paper a look. They discovered that Facebook had not bothered to inform its members about their participation in the experiment, much less ask their consent.

Outrage ensued, as the story pinballed through the media. "If you were still unsure how much contempt Facebook has for its users," declared the technology news site PandoDaily, "this will make everything hideously clear." A *The New York Times* writer accused Facebook of treating people like "lab rats," while *The Washington Post*, in an editorial, criticized the study for "cross[ing] an ethical line." US Senator Mark Warner called on the Federal Trade Commission to investigate the matter, and at least two European governments opened probes. The response from social media was furious. "Get off Facebook," tweeted Erin Kissane, an editor at a software site. "If you work there, quit. They're fucking awful." Writing on Google Plus, the privacy activist Lauren Weinstein wondered whether "Facebook KILLED anyone with their emotion manipulation stunt."

The ethical concerns were justified. Although Facebook, as a private company, is not bound by the informed-consent guidelines of universities and government agencies, its decision to carry out psychological research on people without telling them was reprehensible. It violated the US Department of Health & Human Services' policy for the protection of human research subjects (known as the "Common Rule") as well as the ethics code of the American Psychological Association. Making the transgression all the more inexcusable was the company's failure to exclude minors from the test group. The fact that the manipulation of information was carried out by the researchers' computers rather than by the researchers themselves — a detail that Facebook offered in its defense — was beside the point. As University of Maryland law professor James Grimmelmann observed, psychological manipulation remains psychological manipulation "even when it's carried out automatically."

Still, the intensity of the reaction seemed incommensurate with its object. Once you got past the dubious ethics and the alarming title, the study turned out to be a meager piece of work. Earlier psychological research had suggested that moods, like sneezes, could be contagious. If you hang out with sad people, you may end up feeling a little blue yourself. If you surround yourself with happy folks, your mood may brighten. Kramer and his collaborators (the paper was coauthored by two Cornell scientists) wanted to see if such emotional contagion might also be spread through online social networks. During a week in January 2012, they programmed Facebook's News Feed algorithm — the program that selects which messages to funnel onto a member's home page and which to omit — to make slight adjustments in the "emotional content" of the feeds delivered to a random sample of members. One group of test subjects saw a slightly higher number of "positive" messages than normal, while another group saw slightly more "negative" messages. To categorize messages as positive or negative, the researchers used a standard text-analysis program, called Linguistic Inquiry and Word Count, that spots words expressing emotions in written works. They then evaluated each subject's subsequent Facebook posts to see whether the emotional content of the messages had been influenced by the alterations in the News Feed.

The researchers did discover an influence. People exposed to more negative words went on to use more negative words than would have been expected, while people exposed to more positive words used more of the same — but the effect was vanishingly small, measurable only in a tiny fraction of a percentage point. If the effect had been any more trifling, it would have been undetectable. As Kramer later explained, in a contrite Facebook post, "the actual impact on people in the experiment was the minimal amount to statistically detect it — the result was that people produced an average of one fewer emotional word, per thousand words, over the following week." As contagions go, that's a pretty feeble one. It seems unlikely that any participant in the study suffered the slightest bit of harm. As Kramer admitted, "the research benefits of the paper may not have

justified all of this anxiety."

Big Data, Little People

What was most worrisome about the study lay not in its design or its findings, but in its ordinariness. As Facebook made clear in its official responses to the controversy, Kramer's experiment was just the visible tip of an enormous and otherwise well-concealed iceberg. In an email to the press, a company spokesperson said the PNAS study was part of the continuing research Facebook does to understand "how people respond to different types of content, whether it's positive or negative in tone, news from friends, or information from pages they follow." Sheryl Sandberg, the company's chief operating officer, reinforced that message in a press conference: "This was part of ongoing research companies do to test different products, and that was what it was." The only problem with the study, she went on, was that it "was poorly communicated." A former member of Facebook's Data Science unit, Andrew Ledvina, told *The Wall Street Journal* that the in-house lab operates with few restrictions. "Anyone on that team could run a test," he said. "They're always trying to alter people's behavior."

Businesses have been trying to alter people's behavior for as long as businesses have been around. Marketing departments and advertising agencies are experts at formulating, testing, and disseminating images and words that provoke emotional responses, shape attitudes, and trigger purchases. From the apple-cheeked Ivory Snow baby to the chiseled Marlboro man to the moon-eyed Cialis couple, we have for decades been bombarded by messages intended to influence our feelings. The Facebook study is part of that venerable tradition, a fact that the few brave folks who came forward to defend the experiment often emphasized. "We are being manipulated without our knowledge or consent all the time — by advertisers, marketers, politicians — and we all just accept that as a part of life," argued Duncan Watts, a researcher who studies online behavior for Microsoft. "Marketing as a whole is designed to manipulate emotions," said Nicholas Christakis, a Yale sociologist who has used Facebook data in his own research.

The "everybody does it" excuse is rarely convincing, and in this case it's specious. Thanks to the reach of the internet, the kind of psychological and behavioral testing that Facebook does is different in both scale and kind from the market research of the past. Never before have companies been able to gather such intimate data on people's thoughts and lives, and never before have they been able to so broadly and minutely shape the information that people see. If the Post Office had ever disclosed that it was reading everyone's mail and choosing which letters to deliver and which not to, people would have been apoplectic, yet that is essentially what Facebook has been doing. In formulating the algorithms that run its News Feed and other media services, it molds what its billion-plus members see and then tracks their responses. It uses the resulting data to further adjust its algorithms, and the cycle of experiments begins anew. Because the algorithms are secret, people have no idea which of their buttons are being pushed — or when, or why.

Facebook is hardly unique. Pretty much every internet company performs extensive experiments on its users, trying to figure out, among other things, how to maximize the time they spend using an app or a site, or how to increase the likelihood they will click on an advertisement or a link. Much of this research is innocuous. Google once tested 41 different shades of blue on a web-page toolbar to determine which color would produce the most clicks. But not all of it is innocuous. You don't have to be paranoid to conclude that the PNAS test was far from the most manipulative of the experiments going on behind the scenes at internet companies. You only have to be sensible.

That became strikingly clear, in the midst of the Facebook controversy, when another popular web operation, the matchmaking site OKCupid, disclosed that it routinely conducts psychological research in which it doctors the information it provides to its love-seeking clientele. It has, for instance, done experiments in which it

altered people's profile pictures and descriptions. It has even circulated false "compatibility ratings" to see what happens when ill-matched strangers believe they'll be well-matched couples. OKCupid was not exactly contrite about abusing its customers' trust. "Guess what, everybody," blogged the company's cofounder, Christian Rudder: "if you use the internet, you're the subject of hundreds of experiments at any given time, on every site. That's how websites work."

The problem with manipulation is that it hems us in. It weakens our volition and circumscribes our will, substituting the intentions of others for our own. When efforts to manipulate us are hidden from us, the likelihood that we'll fall victim to them grows. Other than the unusually dim or gullible, most people in the past understood that corporate marketing tactics, from advertisements to celebrity endorsements to package designs, were intended to be manipulative. As long as those tactics were visible, we could evaluate them and resist them — maybe even make jokes about them. That's no longer the case, at least not when it comes to online services. When companies wield moment-by-moment control over the flow of personal correspondence and other intimate or sensitive information, tweaking it in ways that are concealed from us, we're unable to discern, much less evaluate, the manipulative acts. We find ourselves inside a black box.

That suits the interests of the Facebooks and OKCupids of the world, but whether it suits the public interest is a different matter. "How websites work," to pick up on Rudder's cavalier phrase, is not necessarily how they *should* work.

The Right to be Forgotten

Put yourself in the shoes of Mario Costeja González. In 1998, the Spaniard ran into a little financial difficulty. He had defaulted on a debt, and to pay it off he was forced to put some real estate up for auction. The sale was duly noted in the Barcelona newspaper *La Vanguardia*. The matter settled, Costeja González went on with his life as a graphologist. The debt and the auction, as well as the 36-word press notice about them, faded from public memory.

But then, in 2009, nearly a dozen years later, the episode sprang back to life. *La Vanguardia* put its archives online, Google's web-crawling "bot" sniffed out the old article about the auction, the article was automatically added to the search engine's database, and a link to it began popping into prominent view whenever someone in Spain did a search on Costeja's name. Costeja was dismayed. It seemed unfair to have his reputation sullied by an out-of-context report on an old personal problem that had long ago been resolved. Presented without explanation in search results, the article made him look like a deadbeat. He felt, as he would later explain, that his dignity was at stake.

Costeja lodged a formal complaint with the Spanish government's data-protection agency. He asked the regulators to order *La Vanguardia* to remove the article from its website and to order Google to stop linking to the notice in its search results. The agency refused to act on the newspaper request, citing the legality of the article's original publication, but it agreed with Costeja about the unfairness of the Google listing. It told the company to remove the auction story from its results. Appalled, Google appealed the decision, arguing that in listing the story it was merely highlighting information published elsewhere. The dispute quickly made its way to the Court of Justice of the European Union in Luxembourg, where it became known as the "right to be forgotten" case. On May 13 of this year, the high court issued its decision. Siding with Costeja and the Spanish data-protection agency, the justices ruled that Google was obligated to obey the order and remove the *La Vanguardia* piece from its search results. The upshot: European citizens suddenly had the right to get certain unflattering information about them deleted from search engines.

Most Americans, and quite a few Europeans, were flabbergasted by the decision. They saw it not only as

unworkable (how can a global search engine processing some six billion searches a day be expected to evaluate the personal grouses of individuals?), but also as a threat to the free flow of information online. Many accused the court of licensing censorship or even of creating "memory holes" in history.

But the heated reactions, however understandable, were off the mark. They reflected a misinterpretation of the decision. The court had not established a "right to be forgotten." That unfortunate and essentially metaphorical phrase is mentioned only briefly in the ruling, and its attachment to the case has proven a regrettable distraction. In an open society, where freedom of thought and speech are sacrosanct, a right to be forgotten is as untenable as a right to be remembered. What the case was really about was an individual's right not to be systematically misrepresented. But even putting the decision into those more modest terms is misleading. It implies that the court's ruling was broader than it actually was.

The essential issue the justices were called upon to address was how, if at all, a 1995 European Union policy on the processing of personal data, the so-called Data Protection Directive, applied to companies that, like Google, engage in the large-scale aggregation of information online. The directive had been enacted to facilitate the cross-border exchange of data, while also establishing privacy and other protections for citizens. "Whereas data-processing systems are designed to serve man," the policy reads, "they must, whatever the nationality or residence of natural persons, respect their fundamental rights and freedoms, notably the right to privacy, and contribute to economic and social progress, trade expansion and the well-being of individuals." To shield people from abusive or unjust treatment, the directive imposed strict regulations on businesses and other organizations that act as "controllers" of the processing of personal information. It required, among other things, that any data disseminated by such controllers be not only accurate and up-to-date, but fair, relevant, and "not excessive in relation to the purposes for which they are collected and/or further processed." What the directive left unclear was whether companies that aggregated information produced by others fell into the category of controllers. That was what the court had to decide.

Search engines, social networks, and other online aggregators have always presented themselves as playing a neutral and fundamentally passive role when it comes to the processing of information. They're not creating the content they distribute — that's done by publishers in the case of search engines, or by individual members in the case of social networks. Rather, they're simply gathering the information and arranging it in a useful form. This view, tirelessly promoted by Google — and used by the company as a defense in the Costeja case — has been embraced by much of the public. It has become, with little consideration, the default view. When Wikipedia founder Jimmy Wales, in criticizing the European court's decision, said, "Google just helps us to find the things that are online," he was not only mouthing the company line; he was expressing the popular conception of information aggregators.

The court took a different view. Online aggregation is not a neutral act, it ruled, but a transformative one. In collecting, organizing, and ranking information, a search engine is creating something new: a distinctive and influential product that reflects the company's own editorial intentions and judgments, as expressed through its information-processing algorithms. "The processing of personal data carried out in the context of the activity of a search engine can be distinguished from and is additional to that carried out by publishers of websites," the justices wrote.

> Inasmuch as the activity of a search engine is therefore liable to affect significantly […] the fundamental rights to privacy and to the protection of personal data, the operator of the search engine as the person determining the purposes and means of that activity must ensure, within the framework of its responsibilities, powers and capabilities, that the activity meets the requirements of [the Data Protection Directive] in order that the guarantees laid down by the directive may have full effect.

The European court did not pass judgment on the guarantees established by the Data Protection Directive, nor on any other existing or prospective laws or policies pertaining to the processing of personal information. It did not tell society how to assess or regulate the activities of aggregators like Google or Facebook. It did not even offer an opinion as to the process companies or lawmakers should use in deciding which personal information warranted exclusion from search results — an undertaking every bit as thorny as it's been made out to be. What the justices did, with perspicuity and wisdom, was provide us with a way to think rationally about the algorithmic manipulation of digital information and the social responsibilities it entails. The interests of a powerful international company like Google, a company that provides an indispensable service to many people, do not automatically trump the interests of a lone individual. When it comes to the operation of search engines and other information aggregators, fairness is at least as important as expedience.

The court's ruling was not a conclusion; it was an opening. It presented society with a challenge: The mass processing of personal information raises important and complex social, legal, and ethical questions, and it is up to you, the public, to wrestle with those questions and come up with answers.

Our Algorithms, Ourselves

We have had a hard time thinking clearly about companies like Google and Facebook because we have never before had to deal with companies like Google and Facebook. They are something new in the world, and they don't fit neatly into our existing legal and cultural templates. Because they operate at such unimaginable magnitude, carrying out millions of informational transactions every second, we've tended to think of them as vast, faceless, dispassionate computers — as information-processing machines that exist outside the realm of human intention and control. That's a misperception, and a dangerous one.

Modern computers and computer networks enable human judgment to be automated, to be exercised on a vast scale and at a breathtaking pace. But it's still human judgment. Algorithms are constructed by people, and they reflect the interests, biases, and flaws of their makers. As Google's founders themselves pointed out many years ago, an information aggregator operated for commercial gain will inevitably be compromised and should always be treated with suspicion. That is certainly true of a search engine that mediates our intellectual explorations; it is even more true of a social network that mediates our personal associations and conversations.

Because algorithms impose on us the interests and biases of others, we have not only a right, but also an obligation to carefully examine and, when appropriate, judiciously regulate those algorithms. We have a right and an obligation to understand how we, and our information, are being manipulated. To ignore that responsibility, or to shirk it because it raises hard problems, is to grant a small group of people — the kind of people who carried out the Facebook and OKCupid experiments — the power to play with us at their whim.

Works Considered:

Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock, "Experimental Evidence of Massive-Scale Emotional Contagion through Social Networks," PNAS, June 17, 2014.

Christian Rudder, "We Experiment on Human Beings!," OkTrends blog, July 28, 2014.

Court of Justice of the European Union, "Judgment in Case C-131/12, Google Spain SL, Google Inc. v Agencia Española de Protección de Datos, Mario Costeja González," May 13, 2014.

"Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995," *Official Journal of the European Communities,* November 23, 1995.

Sergey Brin and Lawrence Page, “The Anatomy of a Large-Scale Hypertextual Web Search Engine,” *Computer Networks and ISDN Systems,* April 1998.

¤

Nicholas Carr is the author of several influential book on technology and culture, including The Shallows: What the Internet Is Doing to Our Brains*, a 2011 Pulitzer Prize finalist;* The Glass Cage: Automation and US *(2014); and* Utopia is Creepy: And Other Provocations *(2016), reviewed in this collection by Geoff Nunberg.*

3

Algorithms: The Future That Already Happened

by Michael S. Evans (May 31, 2016)

Michael S. Evans critiques Pedro Domingos's The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*, arguing that Domingos's vision of an individually optimized future with a "hypervigilant data citizenry" is unrealistic.*

ONE DAY IN LATE MARCH, Microsoft made a chatbot named Tay. Tay began the day tweeting love to everyone. A few hours later, Tay was quoting Adolf Hitler and offering filthy sex on demand. To borrow a phrase from John Green, Tay fell in love with Hitler and filthy sex the way you fall asleep: slowly, and then all at once. That's because Tay learned from humans, and humans are awful.

Machine-learning algorithms try to make sense of human activity from the data we generate. Usually these algorithms are invisible to us. We see their output as recommendations about what we should do, or about what should be done to us. Netflix suggests your next TV show. Your car reminds you it's time for an oil change. Siri tells you about a nearby restaurant. That loan you wanted? You're approved!

In a sense, you're making these recommendations yourself. Machine-learning algorithms monitor information about what you do, find patterns in that data, and make informed guesses about what you want to do next. Without you, there's no data, and there's nothing for machine learning to learn. But when you provide your data, and when the guesses are correct, machine learning operates invisibly, leaving you to experience life as an endless stream of tiny, satisfying surprises.

Or at least that's how things could be, according to computer scientist Pedro Domingos. In *The Master Algorithm*, Domingos envisions an individually optimized future in which our digital better halves learn everything about us, then go out into the world and act for us, thereby freeing us to be our best non-digital selves. In this vision, machine-learning algorithms replace tedious human activities like online shopping, legal filing, and scientific hypothesis testing. Humans feed data to algorithms, and algorithms produce a better world for humans.

It sounds like science fiction. And it is, notably in Charles Stross's novel *Accelerando*. But is this future possible?

If you're skeptical, maybe it's because you think we're not capable of creating good enough machine-learning

algorithms. Maybe you got a bad Netflix recommendation. Maybe Siri can't understand your instructions. The technology, you might think, just isn't very good.

The Master Algorithm seeks to prove you wrong. Over the course of several chapters on the current state of machine-learning research, Domingos explains that we are close to creating a single, universal learning algorithm that can discover all knowledge, if given enough data. And he should know. In a research field dominated by competition, Domingos has long championed a synthetic approach to machine learning: take working components from competing solutions, find a clever way to connect them, then use the resulting algorithm to solve bigger and harder problems. The algorithms are good enough, or soon will be.

No, the problem isn't the technology. But there are good reasons to be skeptical. The algorithmic future Domingos describes is already here. And frankly, that future is not going very well for most of us.

Take the economy, for example. If Domingos is right, then introducing machine learning into our economic lives should empower each of us to improve our economic standing. All we have to do is feed more data to the machines, and our best choices will be made available to us.

But this has already happened, and economic mobility is actually getting worse. How could this be? It turns out the institutions shaping our economic choices use machine learning to continue shaping our economic choices, but to their benefit, not ours. Giving them more and better data about us merely makes them faster and better at it.

In an article published in *Accounting, Organizations and Society*, sociologists Marion Fourcade and Kieran Healy use the example of credit scoring. Previously you could get credit based on being a member of a large, low-risk group, such as "management employee" or "Ivy League graduate." Individually, you could have a bad year, fail to receive a promotion, or express an unpopular political opinion, yet still access credit that would allow you to, say, invest in real estate, just like the most successful person in your low-risk group.

That is no longer true. Now private corporations ingest ever greater amounts of individual data, including not only our financial transactions, but also our Facebook and Twitter posts, our friends and followers, and our voting histories. These data then feed machine-learning algorithms that classify us based on individual risk profiles, empowering corporations to customize personal financial offerings that extract the most profit from each of us. The good news is you're still able to obtain the same credit as the most successful person in your group. The bad news is that person is you.

More economic options are available than ever before, and to more people. But you'll never see them. And corporations have better information about everyone than we have about ourselves. Rather than improving individual life-chances, machine learning provides more of us with more ways than ever to make our economic lives worse.

So the future, which is the present, isn't looking good for humans. What is to be done?

Domingos's answer is, approximately, learn more about machine learning. *The Master Algorithm* insists on a politics of data in which hypervigilant data citizens actively manipulate the algorithms that might otherwise constrain them. Since machine learning depends on data, Domingos argues, "your job in a world of intelligent machines is to keep making sure they do what you want, both at the input (setting the goals) and at the output (checking that you got what you asked for)."

This ideal of an informed, active citizen probably sounds familiar. It's been shopped around a long time, under

names like "personal responsibility." But this ideal is not reality. In reality, we mostly don't want to know about complex, technical, and consequential processes at all, much less do anything about them.

Take government, for example. In their book *Stealth Democracy*, political scientists John R. Hibbing and Elizabeth Theiss-Morse report that what most Americans want from government has little to do with parties, policies, or influence. What Americans want is government that makes no demands on them whatsoever. The same pattern holds for Wikipedia, the single most-consulted knowledge resource in the world. Millions of people read Wikipedia articles every day. But only a tiny percentage ever writes or edits those articles. Or consider Open Source software, which powers the vast majority of computer servers and no small number of personal computers, phones, and tablets. Anyone with the necessary technical knowledge and programming skill can review, edit, and add to Open Source software. Few ever do.

Even when it results in outcomes that we don't particularly like, most of us want processes that work without our intervention. We just put up with those outcomes. Sure, Wikipedia articles sometimes get the facts wrong. Okay, software bugs crash our computers from time to time. Fine, yes, occasionally an algorithmically driven chatbot becomes a Hitler-quoting sex enthusiast. But how often does that happen, really? Not often enough to do anything about it.

By focusing on individual responsibility, Domingos never fully acknowledges the bigger problem here: even if we wanted to do something, it's not obvious that we could. It's true, as Domingos claims, that machine-learning algorithms can be manipulated if you control the data. Human miscreants trained Tay to be a dirty, dirty chatbot. But then what? Microsoft turned off the algorithm. They can bring it back, or not, and in different form, or not. And they can block every single person who made their bot naughty from ever participating again.

For those in power, machine learning offers all of the benefits of human knowledge without the attendant dangers of organized human resistance. *The Master Algorithm* describes a world in which individuals can be managed faster than they can respond. What does this future really look like? It looks like the world we already have, but without recourse.

The only way that Domingos's vision of the future makes sense is if we press a reset button on the world we have now, eliminating all of the current arrangements and advantages, and starting over. There is precedent. Leviticus 25:8-13 describes a special celebration of liberty called the Year of Jubilee. Every 50 years or so, all property was returned to its original owners, all slaves and prisoners freed, and, at least in theory, the constraints of debts removed. People could start over. They could try again.

So why not a data jubilee? Every few years, delete the data. Start over. Make the machines learn everything anew. Even if we're not sure what's gone wrong, we can at least try turning it off and on again.

¤

Michael S. Evans is a Visiting Scholar in the Department of Mathematics at Dartmouth College. He is the author of Seeking Good Debate: Religion, Science, and Conflict in American Public Life *(University of California Press, 2016)*.

4

Internet Privacy: Stepping Up Our Self-Defense Game

by Evan Selinger (November 10, 2015)

Evan Selinger critically examines Finn Brunton and Helen Nissenbaum's Obfuscation: A User's Guide for Privacy and Protest; *he proposes a multi-pronged digital privacy approach of his own.*

LET'S FACE IT, despite all the coverage "privacy" gets in the post-Snowden world, many of us don't see what all the fuss is about. Or at least we act as if we don't. We avidly connect with friends, family, and colleagues over social media, and for the most part we can't seem to get enough of products and services that use personal information to create customized recommendations: new things to purchase, good songs to listen to, fun places to visit, expedient routes to drive and walk, helpful weather warnings, and pertinent responses to search engine queries. We even like tailored advertising when it's relevant, and we get a tad annoyed when companies aspiring to know us inside and out make poor inferences about our tastes.

Then again, we're also likely to be anxious about the commercial internet and our telltale footprints. Maybe we're dismayed that our past peccadillos or fetishes, our outdated infractions and guilty pleasures won't disappear into the ether. And that's just the tip of the iceberg now that our digital dossiers are overflowing with detail. Few of us happily embrace the idea that every time we go online, companies watch and record our every click — companies that are getting ever better at linking on and offline activity; companies that have their own self-serving agendas and pass off self-interested endeavors as altruistic attempts to make the world a better, more connected place; companies that might be complicit in government surveillance, but only seem embarrassed by their complicity when the public's awareness of that surveillance threatens their revenue.

The fact is that privacy-conscious citizens don't want to be dogged, much less defined, by browsing histories and digitally archived transactions. And they object to third parties ripping their information out of its intended context, as if their intentions and desires are meaningless. Above all, these folks want to be free to reveal different parts of themselves to different audiences without anyone keeping tabs on it all.

Resisting Data Tyranny

When consumers believe corporate tracking leaves them vulnerable, even powerless, they face two basic choices: accept the status quo and remain perpetually anxious about companies being hell bent on reducing our complex, intimate, and ever-changing life stories to reified classifications, or take active steps to limit the encroachment.

In *Obfuscation: A User's Guide for Privacy and Protest*, Finn Brunton and Helene Nissenbaum recommend we step up our self-defense game. They advocate for strong measures, including corporate sabotage.

Brunton and Nissenbaum's outrage centers on two points. First, they insist that major tech companies are actively trying to keep us in the dark, and they're capitalizing on our limited knowledge of how data mining and analysis works. Through corporate secrecy, and by designing mediated environments that don't reveal much about what's going on below the surface (all the while being optimized to nudge maximum user disclosures), as well as by hiring people whose technical knowledge vastly exceeds that of most users, tech companies limit our ability to know what they're currently doing with our information, what they can and might do with it in the near future, and what long-term agendas will kick in as big data repositories grow and information science and technology improves.

Second, Brunton and Nissenbaum maintain that tech companies engage in lots of machinations to exploit our vulnerabilities and keep us powerless. They rig the game so we're inclined to accept without reading terms of service contracts filled with impenetrable legal jargon that we can't grasp and which we couldn't bargain with even if comprehension weren't an issue. They give lip service to "choice," but hugely benefit from the extremely high social and economic costs of opting out of their data-mining services. They also benefit from the fact that private sector attempts to create best practices and codes of conduct are structurally hampered by the "competitive disadvantage associated with general restraints on access to information." And they engage in lobbying efforts that capitalize on the gap between rapid technological advancement and slow legal regulations.

Brunton and Nissenbaum conclude we're living in a state of "data tyranny" that justifies our taking up guerilla tactics.

You might be surprised that Nissenbaum is on the front lines advancing a radical anti-corporate agenda. She's not a stereotypical in-your-face activist, but, like Brunton, a highly respected professor at an expensive private college, New York University. Indeed, Nissenbaum is a prolific author and considered the doyenne of "contextual integrity" — one of the most widely admired and referenced contemporary privacy theories. It spells out what's wrong when a community's sense of appropriate communication and interaction is undermined. Crucially, contextual integrity has institutional teeth, and it directly informs how the Federal Trade Commission approaches privacy protections.

And yet, despite having the attention of influential regulators and advocates, Nissenbaum and Brunton believe the time has come to push for a more radical form of intervention by way of a grassroots movement. It's an admittedly morally complex form of intervention: they recommend that we turn to obfuscatory tactics in order to become deceptive dissenters.

Self-Defense Through Obfuscation

Obfuscation "*is the deliberate addition of ambiguous, confusing, or misleading information to interfere with surveillance and data collection*." There are lots of ways to do it.

Unhappy Facebook users have tried obfuscation by entering vast amounts of false personal information ("Bayesian flooding") to make it harder for Zuckerberg's crew to create robust profiles that advertisers can target. Shoppers who are worried about stores identifying patterns in their spending behavior obfuscate by sharing loyalty cards with one another.

Of course, nothing prevents powerful actors from turning to obfuscation when it suits their purposes. Consider the 2011 Russian deployment of Twitter bots. These "programs purporting to be people" appropriated protestor hashtags, thus adding obfuscatory noise to the communicative system. As a result, dissatisfied citizens had trouble discussing parliamentary elections.

Two technologies exemplify the type of obfuscation Brunton and Nissenbuam are encouraging: TrackMeNot and AdNauseum.

TrackMeNot is an internet browser extension that Nissenbaum collaborated on with Daniel Howe and Vincent Toubiana. It takes your search terms and automatically supplements them with misleading but plausible ones. This makes it more difficult for companies to isolate signal from noise, and allegedly makes it harder for search engine providers, like Google, to discern patterns in queries.

AdNauseum is a Firefox extension for AdBlock Plus that Nissenbaum created with Howe and Mushon Zer-Aviv. On it's own, AdBlock Plus allows users to "block annoying ads, disable tracking and block domains known to spread malware." AdNauseum augments these features by "quietly clicking on all blocked ads while recording, for the user's interest, details about ads that have been served and blocked." Incessant ad clicking makes it harder for trackers to know which clicks are genuine, and it also subverts a contested financial system by paying ad-networks and ad-hosting websites for content that consumers never viewed.

Does Obfuscation Work and Is It Justified?

I have concerns about obfuscation. But I'd also like to take issue with some of the criticisms that have been directed at it.

The obvious practical question is whether TrackmeNot and AdNauseum get the job done. Since their inception, skepticism rightly has been raised about effectiveness.

Efficacy is important, but in my opinion it's not decisive here and certainly shouldn't be used as a trump card to dismiss Brunton and Nissenbaum's agenda. It's early days for their obfuscation efforts, and so we shouldn't be too invested yet in how well the software actually works. Movements take time to build, which is why we're better off viewing TrackMeNot and AdNauseum as akin to thought experiments. They may not leave tech giants shaking, but they do establish "proof of concept," thereby giving coders and their collaborators empirical material to work with: a clear vision that's embodied in a serviceable starting point; a sense of actual functions that can be programmed; and an operational baseline that future projects can improve upon.

Futurist David Brin raises another practical objection, which I'm not buying. He characterizes obfuscation proponents as suffering from "technological myopia" that prevents them from appreciating an ugly truth: their proposals will benefit elites, not average Joes and Janes. Here's what he says:

> Science Fiction author Vernor Vinge referred to this approach [obfuscation] in a novel, portraying a group that filled the Internet with machine-generated garbage information in order to mask personal information beneath a deluge of misleading fabrications. The "Friends of Privacy" thus enabled their own members to conceal their online activities — or so they thought — while making the Net and Web effectively useless at providing reliable information to anyone, anywhere. A classic case of spoiling the commons, in order to achieve a narrow, self-serving goal …
>
> Over the long run, Vernor reveals the obvious — that the "Friends of Privacy" are no more than a front for powerful interests, from oligarchs to criminal gangs, seeking to conceal their nefarious activities

> from any chance of discovery. Indeed, pick any method of concealment or obfuscation — won't elites use it more effectively (by far) than commonfolk? In fact, won't the very programs that you and I purchase, to muddy our footprints, actually report every datum to the original creators of the software? And if you doubt that, oh, have I a bridge to sell you.

Brin is right about one thing. "Commonfolk" probably can't design obfuscation tools on their own and will need help from people who possess sufficient technical expertise. But he's wrong to presume this relationship of dependence is inherently synonymous with corporate dependency, and he should be more imaginative than to speculate that our only option is "purchasing" software.

If open-source advocates are persuaded by Brunton and Nissenbaum's obfuscation arguments, they can in fact build obfuscation tools without surveillance strings attached. Indeed, that's a big reason Brunton and Nissebaum wrote *Obfuscation*. They want to inspire ethically and politically concerned software developers to create new tools that functionally improve upon current ones, and that are easy enough to obtain and use that the charge of "free riding" becomes irrelevant. In other words, they don't want an obfuscating minority-elite exploiting weaknesses in data-tracking systems while the rest of us — the non-techie masses — remain woefully vulnerable. Of course, whether or not elites with deep pockets will benefit disproportionately from having access to expensive obfuscation tools depends on how motivated people and institutions are to providing free high quality user-friendly versions. If such democratized possibilities are ruled out in advance, it's only because the dubious logics of technological and economic determinism are being embraced.

Then, there's the ethical question of whether it's really okay for us to defend ourselves by polluting corporate data pools. Whereas Nissenbaum and Brunton emphasize our limited options for resistance, I think it's important to broaden the frame of analysis. Only then can we discern if more possibilities exist; and only then can we figure out if some of these enable us to avoid ethical murkiness.

Obfuscating isn't the only way to enhance privacy, and while Brunton and Nissenbaum express concern about what information corporations know about us, the fact is that our collective concern about proprietary platforms extends beyond their limited focus. For example, we're apprehensive about the social information we share on such platforms, which becomes available for all kinds of questionable uses. So, while Brunton and Nissenbaum acknowledge they chose the word "obfuscation" "because it connotes obscurity," I think it's helpful to keep in mind the work Woodrow Hartzog and I have done on the subject, which demonstrates that, if your goal is to challenge others who want to access your information, you can in fact select from diverse pathways of resistance. In other words, we should be wary of fixating too intently on any particular approach.

> There are several ways to make online communication more obscure: sharing ideas on platforms that are invisible to search engines; using privacy settings and other access controls; withholding your real name and speaking anonymously or identifying yourself with a pseudonym; disclosing information in coded ways that only a limited audience will grasp; or transmitting content that is encrypted or temporarily accessible through an ephemeral conduit …
>
> Given the prevalence of algorithmic surveillance today, obscurity practices go beyond making it difficult for other people to know what we're saying or that we're the ones saying it. They also include using strategies for conversing online "without tipping […] intentions to the algorithm," effectively "communicating without being computed" … Some of the options to produce obscurity include: referring to folks without tagging them; referring to people and deliberately misspelling their names or alluding to them through contextual clues; sending screenshots of a story instead of directly linking to it …

Thinking about privacy protections through the broader lens of obscurity does more than help us appreciate the diverse ways in which individuals can mitigate both corporate and social surveillance. It also can improve our outlook on the possibilities for genuine policy reform. Federal Trade Commissioner Julie Brill is committed to advancing obscurity protections. She criticizes data brokers and "people search firms" for providing us with inadequate control over our personal information. Approaches like hers serve as positive reminders that obscurity proposals — which can include modes of obfuscation — aren't limited to data guerrilla warfare and the technological arms race that it can create. They can be legal tools for helping ensure that companies promote trust, not fear. Legalizing "a right not to be computed" can, in other words, be part of a multi-pronged approach that has a greater chance of succeeding than Brunton and Nissenbaum suggest.

Books Today

I'd be remiss if I ended this review without addressing the fact that *Obfuscation* is a small and familiar book. The version I read — uncorrected page proofs — only runs 120 pages, and endnotes begin before readers cross the triple page threshold. Moreover, interested readers can get a solid obfuscation education without ever reading the text. They can consult online material that's free and easy to find: Nissenbaum's podcast appearance, "Resisting Data's Tyranny with Obfuscation" (2014), Brunton and Nissenbaum's seminal article, "Vernacular resistance to data collection and analysis: A political theory of obfuscation" (2011), and other quality resources.

In pointing out the size and repetition, I'm not slighting Brunton and Nissenbaum. The fact is, too many academics are long-winded and would benefit from judicious editing of their prose. And as to the re-purposing of ideas, well, scholarship is changing and it's important to acknowledge that this book is a sign of the changing times. Once a researcher is ready to write a game-changing monograph, there's a good chance others have publically commented on the project and the author has already worked out the guiding logic and rehearsed key examples in public: in interviews; in Op-Eds; in blog posts; in podcasts; in archived talks and videotaped conference appearances; in book reviews; and in papers that exist, in some version, outside the confines of paywalls.

Overall, this is a positive development. When you know in advance that the public might scrutinize your claims and when you think hard about the comments your public presentations elicit, you can become crystal clear about your subject matter and develop a reliable sense of how to select examples that arouse your reader's curiosity and sympathy. At the same time, this means turning readable and relevant academic books into unoriginal consolidated records of disparate public archives.

In the case of *Obfuscation*, we shouldn't want it any other way. Movements require public attention.

¤

Evan Selinger is a professor of philosophy at Rochester Institute of Technology, where he is also affiliated with the Center for Media, Arts, Games, Interaction, and Creativity (MAGIC).

5

Algorithmic Life

by Massimo Mazzotti (January 17, 2017)

Using a historical lens, Massimo Mazzotti argues that how we use the word "algorithm" signals a threshold moment — and an opportunity to reconsider existing assumptions about algorithms and their ecologies.

I.

"Algorithm" is a word whose time has come.

Algorithms are changing the worlds we inhabit — and they're changing us. They pop up in op-eds on weighty topics like the future of labor in an increasingly automated world. Writing about how new trading algorithms are taking over Wall Street, a dismayed journalist wonders "which office jobs are next?" — which of us, in other words, will be consigned to the dustbin of irrelevancy? The solution, others gamely counter, may be more algorithms: "How do you find a job? Ask the algorithm." Algorithms promise to bring reliability and objectivity to otherwise uncertain procedures. In 2007, a famous billboard for ASK.com happily capitalized on this promise: it announced to San Franciscans that "the algorithm constantly finds Jesus." Since then, most of us have adjusted our expectations. Algorithms, we have realized, can be carriers of shady interests and vehicles of corporate guile. And so, as a new batch of commentators urge, we must "make algorithms accountable."

Yes, as in the opening of a sci-fi novel, "the algorithms" have arrived. Do they come in peace or are they here to enslave us? This, I argue, is the wrong question.

When we now imagine our future, we think of algorithms. They evoke mixed images of liberation and alienation. To some commentators, algorithms are sublime problem-solvers that will help us end famine and cheat death. They are so effective that their extreme boosters allege they will usher in the end of the human era and the emergence of artificial superintelligence — a hi-tech version of the theological "union with God." For others, the algorithms' promise of power is fraught with ambiguity. In *Homo Deus*, historian Yuval Harari envisages the rise of "Dataism," a new universal faith in the power of algorithms; Dataism will, he contends, help people make sense of a world run by machines just as other religions have helped make sense of earlier dissonant realities. His book is not for readers who think the devil is in the details; it is centered rather on a sweeping Faustian bargain: "humans agree to give up meaning in exchange for power." Growing inequality and the eventual rule of machines might well be part of the price we pay.

Harari's book paints the future as a technological dystopia, which is anything but a new genre. Indeed, the rule of machines is a recurring fear in the modern period. In the 1970s, historian of technology Langdon Winner famously called it the "technics-out-of-control" theme. For him, technology has become largely "autonomous" because humans have ignored its political dimension, and they've done so at their peril. He pointed to the evolution of the very word "technology" from the 19th to the mid-20th century as particularly revealing in this regard; he argued that the meaning of the word had morphed from "something relatively precise, limited and unimportant to something vague, expansive and highly significant," laden with both utopic and dystopic import. The word had become "amorphous in the extreme," a site of semantic confusion — surely a sign, he concluded, that the languages of ordinary life as well as those of the social sciences had "failed to keep pace with the reality that needs to be discussed." The political implications included radicalization of the debate around technology, a proliferation of useless dichotomies, and the disappearance of public spaces for discussing technological change in an informed manner. What has happened to "algorithm" in the past few years, I'd argue, is remarkably similar: we are experiencing a comparable moment of semantic and political inadequacy. But there is more: the term is trying to capture new processes of technological change. How we talk about algorithms can be vague and contradictory, but it's also evocative and revealing. Semantic confusion may in fact signal a threshold moment when it behooves us to revise entrenched assumptions about people and machines.

In what follows, rather than engaging in a taxonomic exercise to norm the usage of the word "algorithm," I'd like to focus on the omnipresent *figure* of the algorithm as an object that refracts collective expectations and anxieties. Let's consider its flexible, ill-defined, and often inconsistent meanings as a resource: a messy map of our increasingly algorithmic life.

II.

As a historian of science, I have been trained to think of algorithms as sets of instructions for solving certain problems — and so as neither glamorous nor threatening. Insert the correct input, follow the instructions, and voilà, the desired output. A typical example would be the mathematical formulas used since antiquity to calculate the position of a celestial body at a given time. In the case of a digital algorithm, the instructions need to be translated into a computer program — they must, in other words, be "mechanizable." Understood in this way — as mechanizable instructions — algorithms were around long before the dawn of electronic computers. Not only were they implemented in mechanical calculating devices, they were used by humans who behaved in machine-like fashion. Indeed, in the pre-digital world, the very term "computer" referred to a human who performed calculations according to precise instructions — like the 200 women trained at the University of Pennsylvania to perform ballistic calculations during WWII. In her classic article "When Computers Were Women," historian Jennifer Light recounts their long-forgotten story, which takes place right before those algorithmic procedures were automated by ENIAC, the first electronic general-purpose computer.

Terse definitions have now disappeared, however. We rarely use the word "algorithm" to refer solely to a set of instructions. Rather, the word now usually signifies a program running on a physical machine — as well as *its effects on other systems*. Algorithms have thus become agents, which is partly why they give rise to so many suggestive metaphors. Algorithms now *do* things. They determine important aspects of our social reality. They generate new forms of subjectivity and new social relationships. They are how a billion-plus people get where they're going. They free us from sorting through multitudes of irrelevant results. They drive cars. They manufacture goods. They decide whether a client is creditworthy. They buy and sell stocks, thus shaping all-powerful financial markets. They can even be creative; indeed, according to engineer and author Christopher Steiner, they have already composed symphonies "as moving as those composed by Beethoven."

Do they, perhaps, do too much? That's certainly the opinion of a slew of popular books on the topic, with titles like *Automate This: How Algorithms Took Over Our Markets, Our Jobs, and the World.*

Algorithms have captured the scholarly imagination every bit as much as the popular one. Academics variously describe them as a new technology, a particular form of decision-making, the incarnation of a new epistemology, the carriers of a new ideology, and even as a veritable modern myth — a way of saying something, a type of speech that naturalizes beliefs and worldviews. In an article published in 2009, entitled "Power Through the Algorithm," sociologist David Beer describes algorithms as expressions of a new rationality and form of social organization. He's onto something fundamental that's worth exploring further: scientific knowledge and machines are never just neutral instruments. They embody, express, and naturalize specific cultures — and shape how we live according to the assumptions and priorities of those cultures.

III.

Historians of science have shown how technological artifacts capture and transform people's imagination, becoming emblematic of certain eras. They are, at once, tools for *doing* and tools for *thinking*. The mechanical clock in Newton's time is a prime example. Consider the heuristic power and momentous implications — scientific, philosophical, and cultural — of seeing the universe as an immense piece of machinery, whose parts relate to one another like those of a sophisticated precision clock. A clockwork universe means, for example, that one can expect to discover regular and immutable laws governing phenomena. But the clock is not simply an image. A clock that can measure seconds and fractions of a second inevitably changes our perception of time. It turns time into something that can be broken down into small units, scientifically measured — and accurately priced. The precision clock helped spawn new temporalities as well as oceanic navigation and the industrial revolution. It was the industrial revolution's metronome. At the same time, the clock was taken to be the best representation of the world it was shaping: a mechanistic, quantifiable, and predictable world, made up of simple elementary components and mechanical forces.

Similarly, seeing the workings of the human mind as analogous to the operations of a hefty Cold War electronic computer signals a momentous cognitive and social shift. Historian of science Lorraine Daston describes it as the transition from Enlightenment reason to Cold War rationality, a form of cognition literally black-boxed in shiny-cased machines, such as the inhumanly perfect monolith in Stanley Kubrick's *2001: A Space Odyssey*. Many sharp minds of the post WWII era believed that the machine's algorithmic procedures, free of emotions and bias, could solve all kinds of problems, including the most urgent ones arising from the confrontation between the two superpowers. It did not work out that way. The world was too complicated to be reduced to game theory, and by the late 1960s the most ambitious dreams of automated problem-solving had been dragged into the mud of the Vietnam War.

Clocks and Cold War computers were, I'm suggesting, *emblematic artifacts*. They shaped how people understood and acted within the world. Clocks and computers also shaped how people understood themselves, and how they imagined their future.

It is in this sense that we now live in the age of the algorithm.

IV.

How accurate is it to say that algorithms *do* things?

The image of the algorithm-as-doer certainly captures a salient feature of our experience: algorithms can alter reality, changing us and the world around us. The Quantified Self movement, which promotes "self knowledge through numbers," is a prime example of how subjectivities can be reshaped algorithmically — in this case, by monitoring vital functions and processing data relative to lifestyles; "the quantified self" engages in data-driven self-disciplining practices.

Algorithms do not just shape subjectivities. The world of our experience is pervaded by code that runs on machines. In the current, expanded sense of the word, algorithms generate infrastructures — like social media — that shape our social interactions. They don't just select information for us, they also define its degree of relevance, how it can be acquired, and how we can participate in a public discussion about it. As media scholar Ganaele Langlois aptly puts it, algorithms have the power to enable and assign "levels of meaningfulness" — thus setting the conditions for our participation in social and political life.

The fact that algorithms create the conditions for our encounter with social reality contrasts starkly with their relative invisibility. Once we become habituated to infrastructures, we are likely to take them for granted. They become transparent, as it were. But there is something distinctive about the invisibility of algorithms. To an unprecedented degree, they are embedded in the world we inhabit. This has to do with their liminal, elusive materiality. In sociological parlance, we could say that algorithms are easily black-boxed, a term I used above to describe how cold war rationality disappeared into computers. To black-box a technology is to turn it into a taken-for-granted component of our life — in other words, to make it seem obvious and unproblematic. The technology is thus shielded from the scrutiny of users and analysts, who cease seeing it as contingent and modifiable, accepting it instead as a natural part of the world. At this point the technology can become the constitutive element of other, more complex, technological systems. Historically, black-boxing has been particularly effective when the technology in question depends on a high degree of logical and mathematical knowledge. Granted, black-boxing does not happen because mathematics is obvious to the general public, but because of the widespread assumption that mathematics consists of deductive knowledge that — as such — is merely instrumental. A technical project like that of a freeway system is, by contrast, saturated with interests; no one would argue for its being economically and politically neutral. But manipulating strings of numbers, or code, according to formal rules? What could possibly be social or indeed biased about that? Aren't these operations purely technical, and therefore neutral?

Not quite. Let me offer an example. Think about the algorithms that produce and certify information. In a programmatic article entitled "The Relevance of Algorithms," media scholar Tarleton Gillespie identified the many ways in which these algorithms have public relevance. As he points out, algorithms select information and assess relevance in very specific ways, and users then modify their practices in response to the algorithms' functioning. Indeed, algorithms produce new "calculated publics" by presenting groups back to themselves. Their deployment is accompanied, observes Gillespie, by "the promise of […] objectivity," whereby "the algorithm is positioned as an assurance of impartiality." These algorithms play a role traditionally assigned to expert groups touting or channeling what might be termed a traditional editorial logic. How did the expert judgment of such groups translate into mechanized procedures? Was this translation straightforward?

Hardly, it turns out. The way algorithms manage information is not simply a mechanized version of that older logic. It is a new logic altogether, an algorithmic logic, which, to quote Gillespie again, "depends on the proceduralized choices of a machine, designed by humans operators to automate some proxy of human judgment or unearth patterns across collected social traces." Endorsing this logic and turning algorithms into trusted information tools is not an obvious or necessary transition: it is a collective *choice* that has important social implications.

If we want to understand the impact of these algorithms on public discourse, concludes Gillespie, it is not

sufficient to know "how they work." We need to examine "why [they] are being looked to as a credible knowledge logic" and which political assumptions condition their dissemination and legitimacy. In other words, we need to be aware of the entanglement of the algorithm with its *ecology* — with the mechanical and human environment within which that particular set of instructions is interpreted and put to work.

To be sure, understanding and visualizing algorithms as embedded in their ecologies is difficult. We tend to default to imagining them in isolation, as self-contained entities. The figure of the algorithm-as-doer reinforces an image of the algorithm as a *tiny machine* crunching data and mixing them up to produce the desired result. You can find its visual rendering in a tongue-in-cheek image of EdgeRank, the Facebook algorithm that decides which stories appear in each user's newsfeed; the image portrays a 19th-century set of three cast-iron grinders, one for each of its main "ingredients": affinity, weight, and time. Such images suggest that algorithms exist independently of their ecology, invariably producing the same effects, wherever and whenever deployed.

That's not what happens with sets of instructions running on physical machines interacting with other systems. But the figure of the algorithm-as-a-tiny-machine is potent. When we look at such an idealized image of a machine, Ludwig Wittgenstein argues in *Philosophical Investigations*, "everything else, that is its movements, seems to be already completely determined." It's a *rigid* and therefore absolutely reliable and predictable object — the stuff of technological dreams, or nightmares.

Experience tells us that real machines do not behave like that. Rather, the idealized machine is a projection of the alleged rigidity of logical rules and mathematical reasoning. All the possible movements of the "machine-as-symbol" are predetermined, writes Wittgenstein; they are already in it "in some mysterious way" — just as, he implies, the correct result of 2 + 2 is already there, as a shadow, when one writes those three symbols.

Summing up, the *deterministic* view of the algorithm — the figure of the algorithm that does things — certainly helps us understand how, as a technological artifact, it can change the world we live in. In this type of speech, the term "algorithm" functions as a synecdoche for software and larger sociotechnical systems. The algorithm-as-doer, however, is also misleading precisely because it hides its larger ecological context; it represents the algorithm as a self-contained mechanism, a tiny portable machine whose inner workings are fixed and whose outcomes are determined. By contrast, an empirical study of algorithms suggests that we can understand their functioning — and their meaning — only by considering the sociotechnical ecologies within which they are embedded.

V.

There is another important reason why the algorithm-as-doer is misleading: it conceals the design process of the algorithm, and therefore the human intentions and material conditions that shaped it.

Thus far, I've argued for the significance of the ecology of algorithms, which is primarily a spatial and synchronic notion. It emphasizes the algorithm's relational properties — how it interacts with machines and human collectives. But we need something more. Consider the example of algorithms that produce and certify information. In exploring their ecology we can address important questions about authority, trust, and reliability. But what about the logic that shaped their design in the first place? Who decided the criteria to be adopted and their relative weight in the decision-making process? Why were the algorithms designed in one particular way and not another? To answer these questions, we need to see the technical features of an algorithm as the outcome of a *process*. In other words, we need a historical — indeed genealogical — understanding of the algorithm. The notion of genealogy is rooted in temporality and diachronicity; it calls attention to the historical emergence of the algorithm's properties, their contingency and precarious stability.

It invites us to question technical features that would otherwise seem obvious and self-explanatory.

A historical sensibility allows us to situate algorithms within a longer quest for mechanization and automation. Like clocks and the programs of early electronic computers before them, current digital algorithms embody an aspiration to mechanize human thought and action in order to make them more efficient and reliable. This is a familiar and yet also unsettling story, constitutive of our modernity. In the most iconic frame from the 1936 comedy *Modern Times*, Charlie Chaplin is literally consumed by the cogs of a mechanized factory. As machines become more efficient, the image suggests, they become more deadly. But what does a term like "efficiency" even mean in this context? Theorizing the antagonistic relationship between labor and machines in 19th-century factories, Karl Marx argued that the new machines not only increased production and reduced costs, but were "the means of enslaving, exploiting, and impoverishing the laborer." Mechanization, he argued, was a weapon in class warfare.

Scholars in the Marxist tradition have continued to pay attention to mechanized production. In *Labor and Monopoly Capital: the Degradation of Work in the Twentieth Century* (1974), for instance, Harry Braverman argued that the kind of automation prevailing in American factories at the time was far from being an obvious imperative. Rather, numerical control machinery had been designed to reshape the relations between management and labor in order to wrest control over the production process from the workshop floor. Other kinds of automation were possible, claimed Braverman, but were not pursued because they would not have had the same *social* effects.

Processes of mechanization and automation, in short, are not simply about productivity and profit. Design features are shaped by social relations. In the 1980s, a wave of social studies built on this intuition to argue that technological artifacts are inherently social — in other words: *any* technology bears the mark of the specific culture and interests that shaped it. The technical sphere is never truly separated from the social sphere: the technical is the social. Technological change is thus never neutral, let alone natural. To say a mechanized procedure is more "efficient" than its predecessor is not an adequate historical explanation for its success. The notion of efficiency is always relative to a set of assumptions and goals. Making these assumptions and goals visible is thus a prerequisite for any informed discussion about technological change and its implications.

How exactly do these insights apply to the study of algorithms? Consider the work of sociologist Donald MacKenzie on the assumptions and negotiations that shaped certain algorithms now embedded into software used by financial traders worldwide. Their design could have been different; there is never just one way to automate a given financial transaction and, more generally, there is never just one way to regulate a market. Choices are made, and these choices do not follow a neutral universal logic; they are the outcome of contingent interests and negotiations. In a similar vein, media scholars Elizabeth van Couvering and Astrid Mager have shown how algorithms behind for-profit search engines are shaped by specific business models, based mainly on user-targeted advertising. They have also shown how these search algorithms stabilize and reinforce the socioeconomic practices they embody.

Precisely because of their unique combination of pervasiveness and invisibility, algorithms can effectively embed and amplify existing social stratifications. The neutrality of algorithms is therefore an illusion. Or, in some cases, a powerful rhetorical tool.

VI.

Digital algorithms should be studied within a long history of mechanization and automation processes. I believe, however, that they also pose new challenges for social scientists. Towards the end of "The Relevance of Algorithms," Gillespie concedes that there might be "something impenetrable about algorithms." Not

only do algorithms work with information on a hard-to-comprehend scale, but they can also be "deliberately obfuscated." In fact, certain "algorithms remain outside our grasp, and they are designed to be." How can we critically scrutinize algorithms if we cannot even grasp them?

The fact that many socially-relevant algorithms lack perspicuity is well known. Choosing to trust them has even been compared to an act of faith — maybe the first step in the rise of Harari's new religion of Dataism. Malte Ziewitz, a science studies scholar, detects a "strange tension" between power and comprehension in the current debate on algorithms. On the one hand, algorithms are continuously invoked as entities that possess precise powers: they judge, regulate, choose, classify, discipline. On the other hand, they are described as strange, elusive, inscrutable entities. It has even been argued, Ziewitz points out, that they are "virtually unstudiable."

What if the algorithms that should set us free of bias are in fact simply better at obscuring how they pull our strings?

Social scientists have decried the difficulties inherent in empirically studying algorithms, especially proprietary ones. This problem is normally framed in terms of "secrecy," a notion that implies strategic concealment. We need, however, a more general concept, like sociologist Jenna Burrell's "opacity."

The reasons why an algorithm is opaque can indeed be multiple and different in kind.

The algorithms considered in these discussions usually use datasets to produce classifications. Burrell's "opacity" refers to the fact that an output of this sort rarely includes a concrete sense of the original dataset, or of how a given classification was crafted. Opacity can be the outcome of a deliberate choice to hide information. Legal scholar Frank Pasquale, a leading critic of algorithmic secrecy, has advocated for new transparency policies, targeting the kind of opacity designed to maintain and exploit an asymmetric distribution of information that benefits powerful social actors.

Algorithms can also be opaque as a result of technical complexity — not in terms of their code per se, but in terms of the structure of the software system within which they are embedded. This point is clearly stated in "Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms," a paper presented in 2014. In this paper, Charles Sandvig and his co-authors point out that "algorithms differ from earlier processes of harmful discrimination (such as mortgage redlining) in a number of crucial ways." One of these is their "complicated packages of computer code crafted jointly by a team of engineers." It follows that "even given the specific details of an algorithm, at the normal level of complexity at which these systems operate *an algorithm cannot be interpreted just by reading it* [emphasis mine]." Grasping how an algorithm actually works requires understanding the entire software structure. Making datasets transparent may be relatively straightforward, but making complex algorithmic processes transparent is not. The ultimate functioning of the algorithm is likely to remain inaccessible to the outsider. The technical choices of certain expert groups can thus have an enormous impact on the collective, but it might be that neither the public nor regulators have the expertise to evaluate their responsibilities and actions.

Still, the kinds of algorithmic opacity described thus far are the outcome of deliberate human decisions. This means that, at least in principle, if the necessary resources and expertise were deployed, it should be possible to understand the inner workings of the resultant algorithms.

In certain cases, however, we have to deal with a more profound kind of opacity, one that does not depend on an information deficit, but rather on the limitations of our cognitive structures.

VII.

This unprecedented type of limitation is a consequence of rapid advances in machine learning. Machine learning is an approach to artificial intelligence that focuses on developing programs that can grow and change when exposed to new data. Typically, it is applied to solve practical problems that cannot be effectively tackled through traditional programming — "by hand," in programming parlance. The programmer's role is thus given to a machine, which adjusts the algorithms based on feedback rather than on traditional logical considerations.

Think of anti-spam software, like the one in your email package. It learns to discern what is spam and what is not, and refines this ability through experience — i.e., your feedback. Now, even in the case of such a simple program, the machine can alter the algorithm to the point that its actual functioning becomes opaque. This means that it might be impossible to explain why the machine classified a particular item as spam or non-spam.

The source of this impossibility is easily located. Human programming follows an overarching logic. For example, the programmer introduces explicit criteria to discriminate between spam and non-spam messages. If the program is modified by a machine, however, the situation changes. First, the perspicuity of the program ceases to be a priority. Its overall structure no longer needs to be surveyable by a human. Second, a machine can increase the scale and the complexity of the decision criteria enormously and — most importantly — it can combine them automatically in ways that do not follow any discernible logic.

Changes in scale and non-linear interactions between features produce the kind of opacity that computer scientists refer to as the problem of interpretability. There is a trade-off, though. Through machine learning, an algorithm can become more accurate at doing what it is supposed to do than an entirely human-designed algorithm could ever be. Choices need to be made. Not just with respect to accuracy versus interpretability, but also with what kind and degree of interpretability we should aim for. Like opacity, transparency is not a univocal term: we need to agree on what exposing the "inner workings" of an algorithm actually means.

The opacity produced by machine learning raises new challenges for those interested in studying how algorithms work and affect our lives. Recent debates on cases of algorithms that have "learned" to discriminate among users based on their racial or gender profile are a concrete example of the kind of issues that social scientists will have to confront. What constraints should be introduced in such algorithms? What is the optimal trade-off between accuracy and "fairness"?

VIII.

Our life is increasingly shaped by algorithmic processes, from the fluctuations of financial markets to facial recognition technology. Manichean arguments for or against digital algorithms are hardly relevant. Rather, we need to understand how algorithms embedded in widespread technologies are reshaping our societies. And we should imagine ways to open them up to public scrutiny — thus grounding shared practices of accountability, trust, and transparency.

This is essential for the simple reason that algorithms are not neutral. They are emblematic artifacts that shape our social interactions and social worlds. They open doors on possible futures. We need to understand their concrete effects — for example, the kinds of social stratification they reinforce. We need to imagine how they *might work* if they were designed and deployed differently, based on different priorities and agendas — and

different visions of what our life should be like.

Algorithms are powerful world-makers. Whose world will they make?

¤

Massimo Mazzotti teaches history of science in the Department of History at UC Berkeley, where he is the Director of the Center for Science, Technology, Medicine, and Society.

6

Opt Out

by Dennis Tenen (February 5, 2016)

Reviewing Bernard E. Harcourt's Exposed: Desire and Disobedience in the Digital Age, *Tenen addresses the surveillance potential of our ever proliferating digital devices; he delivers a call to action.*

LAST SUMMER I spent a few short hours at Rikers Island Correctional Facility with a small group of volunteers from Columbia University. We were there — ostensibly — to run an educational program. Wanting to move away from traditional classroom arrangements, we conceived of our project, Rikers Bot, as a digital storytelling event, combining elements from a programming workshop and a writers group. The idea was to get people excited about programming, while providing a voice to those housed at Rikers under sometimes appalling conditions. We also wanted to give the public at large a glimpse into an otherwise opaque institution. But most importantly, we were there to spend a bit of time with the young people of Rikers: an encounter with the penal system in the midst of our community.

Even in the best of times, entering Rikers is not easy. One goes through a series of gates and checkpoints. One laminated card is exchanged for another. The guard stamps a mark on your wrist using invisible ink.

The jail was not what I expected. Conditioned by the Foucauldian image of the panopticon, I thought that once inside I would observe a state of total surveillance. I expected to see docile bodies, the control of activity, and the swarming of disciplinary mechanisms. Instead, along with many in my group, I was struck by the dilapidated, almost abandoned feeling of the grounds. At one point, something was the matter at the main gate and we were ushered through an alternative entrance. A guard led us through rooms filled with well-used riot gear, guards eating breakfast, and hallways lined with hand-written bills about job safety, potlucks, and charity football games.

We were on our way to a wing of the building that housed the classrooms. But at some point, while walking up the stairs, our guide ducked out into a side door, which locked behind her. We waited in the stairwell. We knew where we had to go, but the doors ahead and behind us were locked. Standing there quietly, together, we understood something about the jail. Despite the guards and the cameras, it was not a place of total surveillance. The building stood in isolation from the outside world. Our presence here was registered in a notebook by hand. We traded our state issued identification for a piece of paper and a stamp. We left no digital trace. In short, no one really knew we were here.

The moment of silence on the stairs punctuated the paradox at the center of Bernard Harcourt's recent book, *Exposed: Desire and Disobedience in the Digital Age*. Real and imaginary panopticons of incarceration from centuries past pale in comparison with those that surround us today. Rather than acquiescing to structures of command and surveillance by force, against our will, and in confinement, we have surrendered to them voluntarily, without duress, and at scale. The condition of willful exposure Harcourt describes in his book challenges well-worn tropes of critical theory. Harcourt begins by noting that we inhabit neither Guy Dubord's "society of spectacle" nor Michel Foucault's "punitive society." Rather, the book documents the emergence of a "new architecture of power relations," "knots of state-like power," and "mini-theaters of consumption."

A new logic of watching and being watched brings with it novel companions: the body double in the reflection of our devices; the quantified self; the ever-vigilant, artificially intelligent personal assistant. Where some wear tracking ankle bracelets in punishment, others strap on tracking wrist bracelets in the name of health and self-discovery. Absent the noise of cell phones and Twitter feeds, the momentary quiet of a stairwell somewhere at Rikers evoked a visceral feeling of isolation. Rikers was scary not because it represented order but because compared to my usually compartmentalized and hyper-connected life it gave space to violence and disorder. "It sounds like I'm in a zoo with all of the birds chirping," wrote one of the young people. "Jail makes us feel deprived and cause vagueness — lacking definite form or character," the bot tweeted. "All the birds outside the windows make it so hard to hear anything."

The expository society, as Harcourt calls this emerging assemblage of technology, practice, norms, and institutions, frustrates long-held intuitions about spectacle and surveillance, inside and outside, public and private. We live in an expository society, Bernard writes, in a society of willful exposure and exhibition. In this perverse light, the inability to expose oneself seems like punishment. And the reward for being watched — liked, favorited, followed — is personal affirmation. Under the emerging regime there is no need for metal bars, cells, or watchtowers. We enter into the hall of mirrors willingly. We demand entrance. And we expose ourselves in return.

To make sense of exposure, Harcourt suggests metaphors that open as yet unexplored critical possibilities. The state is no longer one thing, but many. It is an oligarchic amalgam comprised of the intelligence community, retailers, Silicon Valley, military and corporate interests, social media, politics, banking, and telecommunication. No longer a big brother, it looks rather like an octopus from the badge of the US National Reconnaissance Office, one of the country's 17 intelligence agencies. Instead of the panopticon, we have a pavilion of glass and mirrors. More than an observation tower, the pavilion offers glimpses: it reflects, distorts, and provides for "pockets of obscurity." Like the ivory tower, the panopticons of the past represented remote power. The metaphorical pavilion is close at hand. It is a place to watch and be watched, a place to play and to surveil. Instead of the uncanny body of the impostor double, Harcourt gives us the diligent data double. The data double lives in the spreadsheets of criminal justice statistics and advertising data used to train algorithms. It dutifully embodies our social media habits, our moods, movements, and our "click through" ratios.

We have only begun to understand the personal and political implications of the expository society in which surveillance is both more total and more voluntary than was ever imagined. The nightmare of George Orwell's *1984* is in some ways less intrusive than the reality of 2016. Harcourt's book ultimately points to the desire at the root of our need for exposure. Total surveillance turns out to be pretty enjoyable: Watch people do weird shit online! Share your love for perfectly stacked pancakes with friends! We trade privacy in return for these small bits of fun. Harcourt's work points to the dynamics desire that lead to voluntary exposure. It is a staggering insight: soon, there will be no need to incarcerate. "[W]e will all be watched so closely," he writes. "We won't need to confine anymore, because we will be able to see everywhere" and will be able to "control behavior from a distance."

The banal realities of today's surveillance surpass even the tin foil conspiracies of yesterday. Consider these two devices from 2015 (as a lapsed software engineer among humanists, I continue to regularly follow news and review literature from the tech world). The first is called OnHub made by Google. The marketing materials describe it as a "new type of router for the new way to WiFi." It is supposed to help you stay connected to the internet at home. The device is remarkable for its "innovative circular antenna design," which "provides reliable Wi-Fi coverage in more directions in your home." A tear-down of OnHub by engineers at IFIXIT revealed 13 antennas, along with an ambient light sensor inside.

The second is called Echo, by Amazon. This one is a bit trickier. It is an artificially intelligent personal assistant, similar to Apple's Siri or Microsoft's Cortana, but housed in its own sleek speaker cabinet, meant for the living room. It listens, plays music, and answers questions. "Amazon Echo is designed around your voice," the marketing materials read:

> It's hands-free and always on. […] Tucked under the light ring on Echo is an array of seven microphones. These sensors use beam forming technology to hear you from any direction, even while music is playing. […] Echo is also an expertly tuned speaker that can fill any room with immersive sound […] Echo is built in the cloud, so it is always getting smarter. The more you use Echo, the more it adapts to your speech patterns, vocabulary, and personal preferences.

Uncharacteristically for such gadgets, both OnHub and Echo occupy a central place in the home. Routers and speakers of this sort are usually unsightly, rectangular affairs. They are commonly hidden out of sight, under desks or on top of shelves. The Echo speaker and the OnHub router differ in that they cut an attractive tower-like figure. Both are packed with sensors. The advertisements show them placed on coffee tables and kitchen counters.

Two recent papers in computer science give us an idea of the surveillance potential enabled by the above devices. In "SoundLoc: Acoustic Method for Indoor Localization without Infrastructure" (2014) researchers from UC Berkeley write that: "SoundLoc is a room-level localization system that exploits the intrinsic acoustic properties of individual rooms and obviates the needs for infrastructures." Calling it the "Room Impulse Response measurement system," the researchers develop an algorithm that uses ordinary microphones found on our laptops and cell phones to create a "sound fingerprint" for each room, which allows them to identify the location of its occupants. One would surmise that the seven microphones and the powerful speakers that the Echo carries would make the device particularly well suited for such echolocation, tracking, and fingerprinting of domestic spaces.

A group from MIT published a similar paper titled "See Through Walls with WiFi!" in 2013. "For many years humans have fantasized about X-ray vision," the authors write. Could it be possible today, using wi-fi signals emitted by an ordinary router? As it turns out, yes: "We present Wi-Vi, a wireless technology that uses Wi-Fi signals to detect moving humans behind walls in closed rooms." Just a few years ago the very idea of mapping rooms with microphones or tracking humans behind closed doors with wireless signals would have been relegated to the pages of a novel by Ian Fleming. The suggestion that consumers would voluntarily place such capable devices in the middle of their living rooms still defies belief, despite the evidence. Where a bug in a hotel room caused scandals in the Nixon era, inviting a beastie that carries 13 antennas to the bedroom is now something customers pay $170 for. The fact that the Samsung Smart TV, the United States's best-selling television set, is routinely recording conversations, even when powered off, barely made the news. The terms and conditions policy that comes with the device simply states: "Please be aware that if your spoken words include personal or other sensitive information, that information will be among the data captured and transmitted to a third party through your use of Voice Recognition."

Exposed sounds a timely alarm about the proliferation of such seemingly banal but powerful surveillance mechanisms. If we are to be watched, we must also be able to return the gaze. Forces of capital and control have to grow transparent in proportion to our own transparency. But who will watch the watchers? Acts by whistle-blowers like Chelsea Manning and publishers like Julian Assange suggest a way forward. So do strategies of obfuscation described in a recent volume by Finn Brunton and Helen Nissenbaum. But can these isolated strategies amount to a broader movement for digital resistance? The erosion of privacy and the rise of the expository society happens through countless small, personal decisions, thus the answer must involve a collective response with the participation of us all. The book ends there, calling the reader to action.

But what is actually to be done? Compelled by Harcourt's call to action I suggest the following avenues for exploration in the search for viable models for digital dissent.

First, the war. We must not forget that the state of affairs Harcourt describes in his book is intimately linked with the armaments of perpetual warfare. The expository society is born not out of technological contingencies alone, but also out of the extraordinary measures evoked during wartime. The passage of the USA Patriot Act in the wake of the 9/11 attacks, the Snowden revelations, the Cybersecurity Information Sharing Act tucked into the spending bill passed by Congress at the end of 2015 — these are all connected events the lead to the same source, war. The ill-conceived barter at the heart of all extraordinary measures would trade ideals central to democratic governance for security. War forces our hand. We are told that if we do not comply we will be vulnerable to another terrorist attack. History tells us of a graver danger: tyranny, brought on by rushed measures during seemingly exceptional times. We do not live under a tyrannical regime today. But Harcourt's book does identify infrastructures that have the potential to invite tyranny.

In light of this danger, no less real or pressing than terrorism, I am puzzled by the decline of the antiwar movement. The antiwar rhetoric that defined a generation of conscientious dissenters is no longer a major motivating force for the politically active. According to a recent Gallup poll, Americans worry most about issues like immigration, terrorism, income inequality, drug use, the size and power of federal government, the economy, and climate change. But not war. Violence is perhaps more palatable now because it is hidden and dehumanized by the use of technologies like "precision" strikes and remote drone warfare. But is it possible to restore, in Harcourt's words, "lines between governance, commerce, surveillance, and private life" at home without addressing militarization abroad? I do not think so. Everything near will be exposed in the process of reaching for that hidden and remote threat of terrorism. The perceived difficulty of blowing up an Afghani family on the "deposition matrix" (the US kill list) is used to justify the ease of domestic surveillance. The expository reach reveals all, from here to there. Failing to cease the continual projection of violence abroad, we will be forever unable to stop the erosion of civic society at home.

Second, the economy. The dissolution of the Soviet Union heralded the decline of an ideology: namely statist socialism and central planning. We have yet to deal intellectually with the centralizing forces of capitalist economies. Entities like Google and Amazon exert immense centralizing pressures on the marketplace, including the marketplace of ideas. The technologies Harcourt outlines in his book are all in some way related to the massive aggregation of data. They are both the cause and effect of economies in the process of centralization. Again I ask: Is it possible to address massive surveillance programs without dealing with the underlying monopolies in the marketplace of information? From the engineering systems perspective, the alternative to such central structures is distribution. The emergence of distributed technologies is marked by an attempt to decentralize infrastructure. Witness Diaspora, a federated alternative to Facebook; BitTorrent, a system for distributed file sharing; Bitcoin, peer-to-peer money; Library Genesis, community book collection and sharing; and Wikipedia, a peer-produced internet encyclopedia. These technologies follow a resurgent interest in political thought related to peer production, anarchism, and decentralization. Thinkers like Peter Kropotkin, Leopold Kohr, Lewis Mumford, Jane Jacobs, and Murray Bookchin are making a comeback,

figuring prominently in the conversation about technological resistance. A wider return to this intellectual tradition could fill the critical void left in liberal thought by the decline of communism.

Finally, Harcourt's book makes it impossible not to feel personally complicit in the system by which countless small voluntary transactions, in aggregate, bring the expository society into existence. "Our desires and our passions have enslaved us," Harcourt writes. The emotional affirmation that accompanies exuberant social networking brings with it the very structures of surveillance and control that are used widely by law enforcement and national security agencies. As a community of readers and writers that value intellectual independence, we must begin to align the sharp edge of critical theory with everyday practice. This includes the way we read, write, find, and share information online. For people who spend so much time in front of our computers in search for knowledge, we are remarkably alienated from the material contexts of knowledge production. *Exposed* reminds us of an uncomfortable truth at work in everyday human-computer interaction. We know so little about the smart machines that observe us in our kitchens, offices, and living rooms. Let us cultivate then a keen interest in their literal deconstruction, to learn how they are made and where their wires lead. To strip the aura of fetishism that attaches itself to such magical devices, we need to develop a healthy level of detachment from dumb and inanimate things.

Judging by consumer behavior, the machines that expose us are, for the moment, worth the asking price. They offer enough of a value to overcome the ethical imperative. I have no evidence that Amazon or Google intend to use their hubs and echoes for the purpose of surveillance. There is no need to search for it either, because it is laid out openly, in the terms of service. We are far past the point of caring about the number of cameras or microphones already deployed into action. The idea is not to vilify technology. Technology merely serves our desires.

Harcourt does a masterful job identifying desire as the engine that powers the voluntary surveillance conundrum. Given the chance to better see ourselves in the digital reflection, we will willingly mount many mirrors. The balance between vanity and civic virtue tilts inevitably toward vanity. The burden of ethical reasoning lies ultimately on the part of an informed public. For now, being informed is still less fun than watching smart TV. Consequently, dissent remains the purview of a select few, until the pain of tyranny overwhelms the pleasure of mass exposure.

¤

Dennis Tenen is an Assistant Professor of English and Comparative Literature at Columbia University. He is the author of Plain Text: The Poetics of Computation *(Stanford UP, 2017).*

7

The University of Nowhere: The False Promise of "Disruption"

by Frank Pasquale (November 12, 2015)

Frank Pasquale challenges Kevin Carey's vision of educational disruption in The End of College: Creating the Future of Learning and the University of Everywhere.

FOR TWO DECADES, Harvard Business School professor Clayton Christensen has advanced a sweeping account of "disruption" as an explanation for business history, and as the key to its future. According to disruption theory, nimble competitors replace established firms by developing rival products. Initially cheap and of poor quality, these rival products end up dominating markets. From Amazon to Zillow, disrupters reign.

Christensen's theory of disruptive innovation electrified the consultant class, and its influence soon extended far beyond business. Thought leaders aimed to disrupt government. Christensen co-authored books telling hospital and university leaders to shake up their operations. His public statements now suggest that virtually every facet of human existence can be improved by implementing disruptive principles. Why, he asks, buy a single painting for your apartment, when digital gallerists can email your flatscreen "a fresh piece of art" every three weeks? Disruption has become a theory of everything, set to catapult Christensen to guru status as scholar, consultant, and sage.

And yet the last couple of years have not been kind to him. Historian Jill Lepore's devastating *New Yorker* profile portrayed Christensen as an academic lightweight, who downplays evidence that large, stable companies can sustain their business models. Business researchers Andrew A. King and Baljir Baatartogtokh have strengthened Lepore's case. As Lee Vinsel observes, they found "only 9 of 77 cases that Christensen used as examples of disruptive innovation *actually fit* the criteria of his own theory." Given these embarrassments, it may be time to consign "disruption" to the dustbin of stale management theory buzzwords.

But Christensen's zombie ideas are too politically convenient to disappear — and particularly so in the education sector. Tax-cutting, budget-slashing politicos are always eager to hear that education could be much, much cheaper. The Clayton Christensen Institute had a starring role at a recent Senate Hearing attacking traditional accreditation standards. In Silicon Valley and Wall Street, talk of "disrupting education" mobilizes investors and excites startups. Kevin Carey's *The End of College* is the latest book to seize the imagination of disrupters. It touts massive changes for post-secondary education.

How massive? For Carey, a great deal of instruction should be commoditized, with free or near-free content as

accessible as YouTube videos of massive open online courses (MOOCs). Most research universities shouldn't just shrink. They should "reform" themselves to the point of unrecognizability, or dissolve into the internet ether. We should not mourn them, says Carey, but "shatter" them outright; they are, he believes, "grotesquely expensive and shamefully indifferent to undergraduate learning."
Carey hopes that online courses combined with tiny, impromptu "learning communities" will end college as we know it, replacing it with a "University of Everywhere." His utopian vision, however, is premised on inconsistent values and aims. The likelier result of his policies is a University of Nowhere by way of shifty firms marketing ad hoc vocational education of questionable value or relevance.

The Two Faces of Kevin Carey

Carey's book is an impassioned fusion of prescription and prediction. Following the disrupter's creed, he believes things *need* to change drastically in higher ed, and that they *will* change. But bridging the gap between "is" and "ought" is a formidable task — one Carey tries to solve by muckraking indictments of universities on the one hand and encomia to tech firms on the other. Painting universities in the most unflattering light possible, he sounds a lot like a moral scold, shocking the reader with stories of college students paying "$1,000 for reserved tables and bottles of premium vodka" and faculty decadently pursuing esoteric research agendas. By contrast, he treats the would-be architects of his "University of Everywhere" with kid gloves, praising most of their initiatives with the enthusiasm of a fan-boy.

Neither approach offers a fair account of present university life, nor of the corporations (and corporate-funded foundations) that aim to disrupt it. Rather, they pander to all-too-human predilections for tales of heroism and villainy. I hope to show below that they are particularly jarring when juxtaposed.

As a scold, Carey follows in the footsteps of Ronald Reagan. In the 1960s, the Gipper campaigned for governor of California by promising to "fix the mess at Berkeley." Reagan suspected students were goofing off and smoking pot — and Carey, too, tells college students they need to "put down the bong" and start studying harder. One of Carey's main sources is a study, *Academically Adrift*, which has itself been heavily criticized. The distinguished education researcher Alexander Astin believes that the study's claim "that 45 percent of America's college undergraduates fail to improve their reasoning and writing skills during their first two years of college" simply "cannot be taken seriously." A brutal review of *Academically Adrift* claims that its authors neglect to engage with "hundreds" of articles "showing undergraduate improvement in writing and critical thinking […] although [their book is] aswim with think-tank books and blue-ribbon papers opining the opposite."

Count Carey's *End of College* as one more of those think-tank books, to be brandished in the next iteration of anti-college research, which will then be cited in the next generation of think-tank books. Content cycles are all the rage in disruption circles.

Having painted Hogarthian indictments of university excess and failure early in the book, Carey abruptly embraces the sunny optimism of techno-utopians later on. In *The End of College*, Silicon Valley thought leaders are as pragmatic, nimble, and public-spirited as university administrators are doctrinaire, ossified, and avaricious. They've devised methods of teaching and evaluating students that solve (or will soon solve — Carey vacillates here) all the old problems of distance education.

Online learning at the University of Everywhere could eventually improve outcomes — or degenerate into an uncanny hybrid of *Black Mirror* and *Minority Report.* Big data surveillance will track the work students do, ostensibly in order to customize learning. Get stuck on a lesson? Just keep interfacing with a keyboard, camera, and perhaps haptic sensors. Or perhaps IM some reserve army of tutorial labor via digital labor platforms like

Mechanical Turk or TaskRabbit. Want to prove you aren't faking exams? Just let cameras record your every move and keystroke — perhaps your eye movements and facial expressions, too. According to Carey, "People like [Google's] Peter Norvig will analyze the oceans of information being generated by millions of students and continually optimize and improve what students experience and how much they learn." Certainly we can trust Silicon Valley to respect our privacy and do nothing untoward with the data! "All watched over by machines of loving grace," as the Brautigan poem (and Curtis documentary) put it.

With new platforms like Coursera, Silicon Valley has even lured universities into giving away lectures for free. The colleges think they're establishing good karma with the public, but disrupters hope for a more chaotic endgame: students deciding to watch free courses, then proving their credentials to certifiers who give out "badges" to signify competence in a skill set. The certifiers most likely won't be burdened with any of the teaching, research, community service, counseling (career or otherwise), recreation, social events, extracurriculars, or other long-standing features of residential university communities. They will just verify that student X can do task Y.

It could be a very profitable business. As students pay less for actual instruction by experts, they have more money to spend on badges. Unburdened by legacy staff and faculty, "ed tech" firms could muster a just-in-time workforce to develop new educational technologies. Carey foresees courses produced like movies or videogames — the greatest biology course imaginable, for instance, might cost $100 million to make, but charge a $100 fee to 10 million students, achieving a 10X return for investors. Scale is key, he advises, and therefore "colleges should start thinking about smaller amounts of money multiplied by much larger amounts of people." Meanwhile, investors will continue "unbundling" the university into least-cost providers of content units, student surveillance, and badge-granting.

The Logical Contradictions of Cybernetic Education

It's hard to simultaneously pen a requiem for college, a utopian vision of its low-cost future, and a business plan for educators and investors. Carey blasts universities for operating like for-profit businesses — and then suggests to readers that an entrepreneur who once "settled SEC charges of reporting inaccurate financial results" will treat students better. He praises Great Books maven Robert Maynard Hutchins for criticizing the universities of his day — only to propose a series of nakedly instrumentalist conceptions of education that offer little, if any, room for the classical learning that was among Hutchins's highest concerns.

Sometimes bricolage reflects a healthily hybrid outlook, a willingness to entertain disparate perspectives to achieve a larger vision. But Carey's farrago is dissonant; he grasps at whatever evidence can indict research institutions while ignoring (or gingerly broaching and dismissing) critiques of the "ed tech" he lionizes. Philosophers are worried about regimentation of curricula and homogenization of class content via centralized online courses? Oh, that's just the "philosopher's guild" looking out for itself, a scientist reassures Carey, and the narrative immediately returns to lavish praise for all things tech.

By failing to question that dismissive response, Carey implies that faculty opposition to MOOCs is simply a matter of self-interest. His concerns about greed, so prominent when he discusses universities, fade away when he rhapsodizes about ed tech's "disruptive innovators." Keep in mind that the speculative, financialized business models that Carey claims will transform education have already been tried in other fields, with wildly varying levels of success. We may not care if Pets.com fails spectacularly. But when educational innovation goes wrong, real people suffer. One of Carey's heroes, former Google VP Sebastian Thrun, had a no-bid contract to MOOCify San Jose State University math instruction, only to see the partnership pause after "no more than 51 percent of Udacity students passed any of the three courses" (while "74 percent or more of the students in traditional classes passed"). Eager to find one more excuse to defund higher education, legislators

thrilled to Thrun's siren song of "no frills instruction." Expect *The End of College* to serve a similar function, the "go-to" source for cost-cutters who dismiss research universities as anachronisms.

Silicon Valley is now a magnet for college grads seeking both to "change the world" and get rich. But it can be difficult to square that circle. Carey thinks important ed tech of the future will follow the business model of Google: attract millions of customers, and somehow monetize that audience later. He never pauses to reflect on the legal and ethical controversies dogging that firm, ranging from antitrust to privacy violations. Nor does he betray any serious worries about rapid centralization and reuse of student data by under-regulated firms, or the vast critical literature on Silicon Valley solutionism, or black boxed instructional technology run by algorithms that can't be accessed by the students it is assessing.

Finding Resources for a Better Future

Two reviewers with deep knowledge of educational technology and policy have already dismissed as "techno-fantasies" Carey's scorched earth dreams of corporate "thunder lizards" (his words, not mine) disrupting universities. Traditional college education endures — and even those who dismiss it rarely, if ever, try to dissuade their own children from attending a university. If colleges were really so terrible at preparing the workforce, the college earnings premium would have disappeared long ago. In addition, most employers are unlikely to subscribe to Carey's biased bashing of colleges. The alternatives he is pushing are too untested and implausible.

So why bother reading Carey? Because, like Donald Trump blustering his way to the top of the Republican field by popping off shocking sentences, Carey's rhetoric has political heft. To the extent it gains traction among education policy staffers (and the student loan companies that love to hire them), it changes the debate. *The End of College* is a master class in translating an elite project of privatization and austerity into bite-sized populist insults, even as it sings the praises of powerful corporations.

Carey claims he wants dramatically better educational opportunities for all. But that goal will require *more* public support — the kind of budget allocations that are hard to justify when a think tanker opines that Silicon Valley models (for providing free information in exchange for personal data) are on the cusp of disrupting higher ed. Finding new balances between human-led and software-driven classroom exercises (as well as lectures, tests, capstone projects, essays, and exams) demands experimentation and pilot programs. Productive deployment of new learning models depends on *humane* and *distributed* automation that balances the insights of diverse domain experts (including current educators) and technologists. Congress and states should help colleges pay adjuncts and education technologists fairly, while reducing student debt burdens.

Bernie Sanders has proposed a plan that would aid colleges in such a digital transition by directing funds away from sports and toward "activities that improve instructional quality and academic outcomes." But Carey has mocked Sanders's plan as quaint, a "kind of weird" idea from a "rumpled, redistributionist" utopian — even as he celebrates Silicon Valley utopianism. He has called educational policy analyst Diane Ravitch's style "increasingly dismissive and strident," concluding that "her righteousness can be breathtaking." But *The End of College* is itself "dismissive and strident" in its condemnation of universities that are often struggling to balance competing goals. It self-righteously alleges that research universities have grown by exerting political influence and by using accreditation to knock out competitors. Meanwhile, Carey fails to reckon with the untoward forces that might be driving his own book's popularity among elites.

The End of College will undoubtedly please many donors to, and board members of, the New America Foundation — its author's employer. Many millionaires and billionaires want to see their taxes go down — and what better way to accomplish that than to tell Congress and state legislatures that college courses can and

should be no more costly than a video game? A whole new round of budget cutting can ensue. After reviewing the funding and history of the New America Foundation, is it hard to imagine why Carey would push Google and tech firms as the saviors of education? Carey himself points out the symbiosis between his "University of Everywhere" and tech business interests, emphasizing how eager they are to grab a slice of the $4.6 trillion global education market.

In short, if Carey is going to portray so many people working in higher education as self-serving and grasping, he should not be surprised when critics see his work in the same light. And when an author's critical faculties so often vanish when he discusses the technology sector that supports his employer, this raises troubling questions about think-tank expertise in general. Before touting D.C. researchers' "findings" and "big idea books," the media and indeed all of us should look closely at exactly what interests are funding the think tanks behind them. There probably won't be much room on the curriculum at the "University of Everywhere" for such inquiry — which makes it all the more vital for it to happen now. Researchers like Audrey Watters are developing a more humane vision for education technology. Let's hope policymakers begin treating their careful and rigorous work with a level of enthusiasm and interest now reserved for big foundations and billionaires.

¤

Frank Pasquale is an expert on the law of big data, predictive analytics, artificial intelligence, and algorithms, and has served as a member of NSF-sponsored Council on Big Data, Ethics, and Society. His book The Black Box Society: The Secret Algorithms That Control Money and Information *(Harvard, 2015) developed a social theory of reputation, search, and finance.*

8

Myth, Magic, and Mysticism in the Age of Information

by Erik Davis (March 29, 2015)

Erik Davis's TechGnosis *is considered the classic text on the relations between technology and the occult. This "Afterword 2.0" was written for a new North Atlantic Books edition.*

IT MAKES ME SLIGHTLY PAINED to admit it, but the most vital and imaginative period of culture that I've yet enjoyed unfolded in the early 1990s (with the last years of the 1980s thrown in for good measure). There was a peculiar feeling in the air those days, at least in my neck of the woods, an ambient sense of arcane possibility, cultural mutation, and delirious threat that, though it may have only reflected my youth, seemed to presage more epochal changes to come. Recalling that vibe right now reminds me of the peculiar spell that fell across me and my crew during the brief reign of David Lynch's *Twin Peaks*, which began broadcasting on ABC in the spring of 1990. Today, in our era of torrents, YouTube, and Tivo, it is difficult to recall the hold that network television once had on the cultural conversation, let alone the concrete sense of historical time. Lynch's darkside soap opera temporarily undermined that simulacra of psychological and social stability. Plunging down Lynch's ominous apple-pie rabbit hole every week, we caught astral glimmers of the surreal disruptions on the horizon ahead. I was already working as a culture critic for the *Village Voice*, covering music, technology, and TV, and later that year I wrote an article in which I claimed that, in addition to dissolving the concentrated power of mass media outlets like ABC, the onrushing proliferation of digital content channels and interactive media was going to savage "consensus reality" as well. It wasn't just the technology that was going to change; the mass mind itself was, in an au courant bit of jargon from Deleuze and Guattari's *A Thousand Plateaus*, going molecular.

Molecular meant a thousand subcultures. Pockets of alternative practices across the spectrum crackled with millennialist intensity in the early nineties, as if achieving a kind of escape velocity. Underground currents of electronic music, psychedelia, rap, ufology, cyberculture, paganism, industrial postpunk, performance art, conspiracy theory, fringe science, mock religion, and other more or less conscious reality hacks invaded the spaces of novelty and possibility that emerged in the cracks of the changing media. Hip-hop transformed the cut-up into a general metaphor for the mixing and splicing of cultural "memes" — a concept first floated by Richard Dawkins in 1989. Postmodernism slipped into newsprint, Burning Man moved to the desert, and raves jumped the pond, intensifying the subliminal futurism of American electronic dance music into a sci-fi hedonism that turned the DJ into a nerdy shaman and the nervous system into a launching pad. The ambient music designed to fill chill-out tents helped stage a return of a pop-tech mysticism, intensified by MDMA's glowing body-without-organs and the return to serious psychedelia aided and abetted by Terence McKenna

and other Internet-enabled psychonauts. The eighties zine scene continued to flourish, but new production tools allowed publications like *Mondo 2000, Magical Blend, Gnosis*, and the "neurozine" *Boing Boing* to catapult from the DIY zone onto the magazine racks. At the same time — and with enormous effect on the weirdness to come — the zine ecology began colonizing the online hinterlands of BBSes, Usenet alt groups, and the Well (which was, well, a big BBS). Even cable access TV was getting pretty strange (at least in Brooklyn). Some wags joked that Hendrix had rightly prophesied, and that the sixties had turned out to be the nineties after all. And while that fantasy radically distorted the street politics of the former and the technology-primed economics of the latter, it did announce that the old hippie divide between a computerized technocracy and an earthy analog underground had not only broken down but dissolved.

This was, quite simply, an awesome time to be a cultural critic. At the *Village Voice*, then a feisty paragon of identity politics and primo alternative journalism, I was encouraged by a handful of highly skilled (and highly tolerant) editors to write about everything from cosmic heavy metal to posthumanist philosophy to *The X-Files* to the Zippies. Following the steps of my *Voice* pal and fellow tech journalist Julian Dibbell, I got a Panix dial-up account in 1993, and dove into the weirdness of alt groups, French theory listservs, and the social experiments of LambdaMOO, where I encountered a crew of highly intelligent and creative anarchist pagans that blew my mind. Those years were, by far, the most fun I ever had online. But the real initiation into the stream of technomagic that inspired *TechGnosis* occurred a couple years earlier, when I flew from New York to the Bay Area in order to attend the first and only Cyberthon, a paisley-flaked technology gathering whose speakers included Timothy Leary, Terence McKenna, and Bruce Sterling. Virtual reality, now making a belated comeback through Oculus Rift and related gamer gear, was all the rage. I strapped on dread-headed Jaron Lanier's data-glove rig, and I toured the VR lab at NASA Ames with the deeply entertaining John Perry Barlow. I met a sardonic William Gibson, who single-handedly engineered our "collective hallucination" of cyberspace, and a standoffish Stewart Brand, whose *Whole Earth* tool fetishism presaged the Cyberthon's meet-up of counterculture and cyberculture. For me, born in the Bay Area but raised and living on the East Coast, the Cyberthon was a strange kind of homecoming: one that only strapped me onto a new line of flight, a cruise that rode the growing updrafts of what would become the mass digital bloom.

TechGnosis was in many ways woven from the travel diary of that cruise. As a journalist, as well as a heady seeker of sorts, I was already devoted to tracking the juxtaposition of spirituality and the material grit of popular culture, a juxtaposition that in the nineties came to include new technologies, human augmentation tech, and the dawning "space" of digital mediation. Once I tuned into this techgnostic frequency, I realized that the waves radiated backward as well as forward, not just toward Teilhard's apocalyptic Omega Point or McKenna's jungle Eschaton, but toward the earliest technical stirrings of Paleolithic Homo sapiens. I became seized by the McLuhanesque conviction that the history of religion was really just a part of the history of media. As a pagan dabbler, I grokked that the hermetic and magical fabulations that had gone underground in the modern West had returned, like Freud's repressed hankerings, in technological forms both built and imagined, demonic and transcendent, sublime and ridiculous. I began to track these secret histories, and my notes grew until they demanded to be a book.

Today there is so much wonderful and intelligent material on occult spirituality — in scholarship, literature, and the arts — that it is hard to remember just how esoteric this stuff was in the nineties. Peers at the time suggested that, outside certain recondite circles, my research might prove bootless given the more pressing issues — and pragmatic opportunities — associated with the digital revolution. And yet, as the pieces fell into place, as I befriended technopagans or stumbled across cyborg passages in hermetic texts, I felt I no longer had choice in the matter. I was possessed by what Teilhard had called the "demon (or angel) of Research," which is one way of describing what takes place when the object of study turns around and grabs you by the balls. I had to write *TechGnosis*. And though other writers and historians were tuned into these questions both before and alongside me, I am chuffed, as the British say, that scholars, hackers, mystics, and artists alike

continue to draw from the particular *Wunderkammer* I assembled.

I think *TechGnosis* continues to speak despite its sometime anachronism because it taps the enigmatic currents of fantasy, hope, and fear that continue to charge our tools, and that speak even more deeply to the profound and peculiar ways those tools shape us in return. These mythic currents are as real as desire, as real as dream; nor do they simply dissipate when we recognize their sway. Nonetheless, technoscience continues to propagate the Enlightenment myth of a rational and calculated life without myths, and to promote values like efficiency, productivity, entrepreneurial self-interest, and the absolute adherence to reductionist explanations for all phenomena. All these day-lit values undergird the global secularism that forms the unspoken framework for public and professional discourse, for the "worldview" of our faltering West. At the same time, however, media and technology unleash a phantasmagoric nightscape of identity crises, alternate realities, memetic infection, dread, lust, and the specter of invisible (if not diabolical) agents of surveillance and control. That these two worlds of day and night are actually one matrix remains our central mystery: a rational world of paradoxically deep weirdness where, as in some dying earth genre scenario, technology and mystery lie side-by-side, not so much as explanations of the world but as experiences of the world.

Take the incipient Internet of things — the invasion of cheap sensors, chips, and wirelessly chattering mobile media into the objects in our everyday world. The nineties vision of "cyberspace" that partly inspired *TechGnosis* suggested that a surreal digital otherworld lay on the far side of the looking glass screen from the meatspace we physically inhabit. But that topology is being decisively eroded by the distribution of algorithms, sensing, and communicating capabilities through addressable objects, material things that in some cases are growing extraordinarily autonomous. There are sound reasons for these developments, which arguably will greatly increase the efficiency and power of individuals and organizations to monitor, regulate, and respond to a world spinning out of control. As such, the Internet of things offers consumers another Gernsback carrot, another vision of a future world where desire is instantly and transparently satisfied, where labor is offloaded onto servitors, and where we are all safely watched over by machines of love and grace. But if the social history of technology provides any insight at all — and I would not have written *TechGnosis* if it didn't — this fantasy is necessarily coupled to its own shadow side. As in the tale of the sorcerer's apprentice, algorithmic agents will be understood as possessing a mind of their own, or serve as proxies for invisible agents of crime or all-watching control. Phil Dick's prophecy, cited earlier in *TechGnosis*, is here: our engineered world is "beginning to possess what the primitive sees in his environment: animation." In other words, a kind of anxious animism, the mindframe once (wrongly) associated with the primitive origins of religion, is returning in a digitally remastered form. Intelligent objects, drones, robots, and deeply interactive devices are multiplying the nonhuman agents with whom we will be forced to negotiate, anticipate, and dodge in order to live our lives. Sometimes remote humans will be at the helm of these artifacts, though we may not always know whether or not people are directly in the loop. But all of it — the now wireless world itself — will become data for the taking. So if Snowden's NSA revelations felt like the cold shadows of some high-flying *nazgûl* falling across your backyard garden, get ready to be swallowed up in the depths of the uncanny valley.

One side of this new animism we already know by another name: paranoia, which will continue to remain an attractive (and arguably rational) existential option in our networked and increasingly manipulated world. Even if you set aside the all-too-real problems of political and corporate conspiracy, the root conditions of our hypermediated existence breed "conspiracy theory." We live in an incredibly complicated world of reverberating feedback loops, one that is increasingly massaged by invisible algorithmic controls, behavioral economics, massive corporate and government surveillance, superwealthy agendas, and insights from half a century of mind-control ops. It is impossible to know all the details and agendas of these invisible agents, so if we try to map their operations beyond the managed surface of common sense and "business-as-usual," then we almost inevitably need to tap the imagination, with its shifty associative logic, as we build our maps and models out of such fragmentary knowledge. That's why the intertwingled complexities — aided and abetted

by the myopic and self-reinforcing conditions of the Internet — found even in the most concrete conspiracy investigations inevitably drift, as systems of discourse, towards more arcane possibilities. The networks of influence and control we construct are fabulated along a spectrum of possibility whose more extreme and dreamlike ends are effectively indistinguishable from the religious or occult imagination. *JFK = UFO*. Analyses of the "twilight language" hidden in the latest school shooting, or Illuminati hand signs in hip-hop videos, or the evidence for false flag operations buried in the nitty-gritty data glitches of major news events — all these disturbing and popular practices suggest an allegorical art of interpretation that is impossible to extricate from our new baroque reality, with all its reverberating folds of surface and depth. Paranoia's networks of hidden cause not only resonate with the electronic networks that increasingly complicate and characterize our world, but suggest the ultimate Discordian twist in the plot: that the greatest forms of control are the stories we tell ourselves about control.

Indeed, the most obvious place to track the prints of myth, magic, and mysticism through contemporary technoculture is, of course, in our fictions. At the beginning of the nineties, geek culture was largely a nerdy niche, its genres and fannish behaviors leagues away from serious cool. But as geeks gained status in the emerging digital economy, the revenge of the nerds was on. The battle is now over, and the nerds rule: popular culture is dominated by superheroes, science fiction, sword and sorcery, RPGs, fanfic, Comicons, Lovecraftmania, cosplay. Geek fandoms have gone thoroughly mainstream, propagated through gaming, Hollywood, online newsfeeds, massive advertising campaigns, and office cubicle decor. With a qualified exception for hard SF, these genres and practices are all interwoven, sometimes ironically, with the sort of occult or otherworldly enchantments tracked in *TechGnosis*. But its not just geek tastes that rule — it's geek style. As the software analytics company New Relic put it in a recent ad campaign, we are all "data nerds" now. In other words, we like to nerd out on culture that we increasingly experience as data to play with. The in-jokes, scuttlebutt, mash-ups, and lore-obsession of geekery allow us, therefore, to snuggle up to the uncanny possibilities of magic, superpowers, and cosmic evil without ever losing the cover story that makes these pleasures possible for modern folks: that our entertainments are "just fictions," diversions with no ontological or real psychological upshot, just moves in a game.

The funny thing about games and fictions is that they have a weird way of bleeding into reality. Whatever else it is, the world that humans experience is animated with narratives, rituals, and roles that organize psychological experience, social relations, and our imaginative grasp of the material cosmos. The world, then, is in many ways a webwork of fictions, or, better yet, of stories. The contemporary urge to "gamify" our social and technological interactions is, in this sense, simply an extension of the existing games of subculture, of folklore, even of belief. This is the secret truth of the history of religions: not that religions are "nothing more" than fictions, crafted out of sociobiological need or wielded by evil priests to control ignorant populations, but that human reality possesses an inherently fictional or fantastic dimension whose "game engine" can — and will — be organized along variously visionary, banal, and sinister lines. Part of our obsession with counterfactual genres like sci-fi or fantasy is not that they offer escape from reality — most of these genres are glum or dystopian a lot of the time anyway — but because, in reflecting the "as if" character of the world, they are actually realer than they appear. That's why we have seen the emergence of what scholars call "postmodern religion" between the cracks of our fandoms: emotionally wrenching funerals on *World of Warcraft*, Mormon (and Scientological) science fictions, Jedi Zen, even Flying Spaghetti Monster parodies that find themselves wrestling with legal definitions of "real" religion.

But it may be in horror that we most clearly see the traces of technological enchantment today, a trace as easy to track as the eerie frame of Slender Man. Emerging from the mines of creepypasta, a hard-geek zone of Internet-enabled horror tales designed to propagate virally, Slender Man first appeared as a faceless and abnormally tall spook in a black suit lurking in the background of an otherwise placid playground scene posted to the comedy prankster site *Something Awful*. Memetically, Slender Man had the goods, and soon

found himself multiplied through a vast number of images, videos, cosplay costumes, and online narratives. I like to think Slender Man's popularity may have had something to do with his resemblance to the lanky and reserved H. P. Lovecraft — a resonance underscored by the tentacles he sometimes sports. Lovecraft's so-called Cthulhu mythos is the paragon of that weird interzone aimed at by so many horror franchises, which seek to achieve an "as if" reality through self-referential and intertextual play that seems to bring the phantasm further into being. This play, it could be argued, is almost what the Internet is designed for. But here we speak not of fell Cthulhu, nor of the dreaded *Necronomicon* and its various incarnations in print. Instead, it was the gangly Slender Man who stepped from cyberspace into the real when two twelve-year-old girls from Wisconsin — perhaps not unlike the adolescents in the original Photoshopped playground image — stabbed a classmate in the woods in order to please the crowd-sourced wraith. The possible mental instability of the girls is not really the point here — it is the rapid Net-enabled mediation of fictions into something more like folklore, but a folklore now rendered viral and invasive through the virtual and social media that increasingly circulate and condition "consensus reality." Less horrifying examples of this sort of phantasmic logic can also be found in the fringe phenomena of Otherkin and tulpamancy — Internet-fueled subcultures that proclaim the ontological reality of beings and identities cobbled together in part from fandom and modern folklore, but gaining their consistency through the digital mediation and collective construction of unusual psychological experience.

In a recent essay for the book *Excommunication*, Eugene Thacker examines the constitutive role that media have played in many supernatural horror tales. In normal life, the different times and places that communication technologies tie together belong to the same plane of reality — New Caledonia may be an exotic place, but when I FaceTime someone there, I am still communicating with a locus in Terran spacetime. But in supernatural horror, media create portals between different orders of reality, what Bruno Latour would call different ontological "modes." Examples include the cursed videotape in the J-horror classic *Ringu*, or the device in Lovecraft's "From Beyond" that reveals the normally invisible beasties that flit about our dimension. The paradox of such fictions is that the remoteness of the otherworld is made immanent in the technology itself, present to hand in an actual artifact that still oozes otherness. The device it grows haunted, or weird, not because the technology breaks down, but because it works too well. Glitches, noise, and stray signals are no longer technical faults but the flip side of another order of being leaking through. Though Thacker is interested in horror fiction, a similar bleed between ontological realms occurs in some paranormal practices. Take the legions of photographers drawn to angels, ghosts, and manifestations of the Blessed Virgin Mary. Though the ubiquity of phones and post-processing techniques should, according to the rationalist rules of evidence, diminish the believability of specters or heavenly beings, some photographers have developed a rich iconography of lens flares, floating orbs, streakers, and other mysterious marks that indicate otherness. Media will always present technical anomalies, and such anomalies will always offer stages for oracular and otherworldly perception, whether or not you consider such perceptions as internally-generated apophenic projections, or as living traces of those mysterious orders of presences that seem to ghost communication.

The spaces of novelty that *TechGnosis* explored were largely opened up by developments in technical media, including the digital revolution that emerged at the end of the twentieth century. But a far more fundamental example of Thacker's "weird media" remains the human sensorium itself: at once the realtime flux of perception, feeling, and cognition, as well as the neural substrate that conditions, and arguably causes, this ongoing mediation of reality. And it is the human sensorium, conscious and unconscious, that has become the ultimate object of technical manipulation, augmentation, and control. In part, this represents the steady march of technoscience and the rational Enlightenment project it represents, and as such would seem to suggest that we are close to banishing all those hoary ghosts of yore. But there is a funny paradox about the neuroscientific bid to map the workings of the mind: the more totalizing the effort to explain consciousness and all its features, the more seriously researchers must engage, in a non-trivial manner, the most marvelous and otherworldly events: lucid dreams, placebo healings, out-of-body journeys, near-death experiences, extreme-

sports highs, meditative insights, DMT otherworlds, and a whole host of apparitions, premonitions, and other paranormal phenomena. While intricate (and intransigent) sociobiological explanations for all this weirdness will continue to be presented as the only serious game in town, and while the organized (and well-funded) armies of militant skeptics will continue to fan the smokescreen that surrounds serious parapsychological research, the phenomena themselves must be taken seriously as experiential realities. Weirdness, in other words, cannot simply by swept under the rationalist carpet — it is thoroughly woven into the world that needs to be explained — and that will continue to be experienced, above and beyond all explanation.

In "The Spiritual Cyborg," for my money the most important chapter in this book, I talked about how the extreme view of human being presented by reductionist science — that we are basically neo-Darwinian DNA robots — has itself been hijacked by some techgnostics for the purposes of mystic liberation and visionary reality programming. This unexpected twist, by which skepticism becomes a tool of spirit, is one of the key points of the anthropological perspective I favor. It is not that religious visions or spiritual values or occult cosmologies are special, unvarnished forms of truth. They are indeed stories and constructions, fabulations and fabrications we use (and mis-use) to get by. The point instead is that the demands of cold hard reality — whether those are framed as reductive naturalism, economic pragmatism, or a harsh and arrogant skepticism that does injustice to all manner of realities hard and soft — are also stories and constructions. Facts are very special objects, which is why they must be constructed through such careful and painstaking methods. But they are still human fabrications, especially when we noise them abroad through popular media or glue them together into more or less impervious worldviews. We are all wearing tool belts here, scientists and mystics alike, fashioning experience into artifacts and realities that feedback on us, inevitably, as stories, shaping us in turn. Backed up by sociobiology or neuroscience — or the pop-science simulacra of sociobiology and neuroscience — many of today's dominant technological stories are devoted to augmenting the competitive advantage of the same old rational agent, or, more insidiously, to manipulating subjectivity for purposes of economic or social control. Instead, I hope that we rapidly and creatively expand our range of what the German philosopher Peter Sloterdijk calls "anthropotechnics" — those processes and practices that turn us into perceiving subjects, that train our capacities, that bootstrap our own transformation. Rational calculation should never tame what Sloterdijk describes as the "vertical tension" that pulls us ever upward and outward, and toward the acrobatics of the spirit.

I admit that sometimes this seems like a thin hope indeed. A massage therapist I know up in Northern California recently remarked that, faced with the apparent gloaming of human history, most folks she knew were either rooting themselves in more embodied, local, and offline lives, or were diving with more mutant gusto into the intertwingled webwork of the digital *cosmopolis*. I have always been a fan of the "middle way" — between reason and mystery, skepticism and sympathy, cool observation and participation mystique. Facing the technological future, I remain a being of ambivalence, suspended, like many, I suspect, in a vexed limbo of bafflement, wonder, denial and despair. I remain fascinated and amazed by our realtime science fiction and the cognitive enhancements (and estrangements) provided by our increasingly posthuman existence. But I also find myself profoundly alienated by the culture of consumer technology, aggravated by the fatuous and self-serving rhetoric of Silicon Valley tools, horrified by our corporatized surveillance state, and saddened by the steely self-promoting brands that so many people, aided and abetted by social media, have become. I was born in the Summer of Love, and while my generation had the uncanny privilege of witnessing the dawn of mass digital culture, I increasingly find myself communing with the other side of the coin: the analog sunset it has also been our blessing to witness and undergo. Like the warm crackles of vinyl, or the cosmic squiggles of a wild modular synth, or the evocative glow of an actual Polaroid, the resonant frequencies of a less networked world still illuminate all my relations. I do not feast on nostalgia, but nostalgia is not the same thing as affirming the gone world that still signals us now, in the timeless time of transmission.

¤

Erik Davis is the author of TechGnosis: Myth, Magic, and Mysticism in the Age of Information*,* Led Zeppelin IV *(for the 33 1/3 series),* The Visionary State: A Journey through California's Spiritual Landscape*, and* Nomad Codes: Adventures in Modern Esoterica*.*

9

Creepy Futures: Nicholas Carr's History of the Future

by Geoff Nunberg (November 21, 2016)

Reviewing Nicholas Carr's Utopia is Creepy: And Other Provocations*, a collection of Carr's writings from 2005 to 2015, Geoff Nunberg argues that this era's techno-utopianism is uncannily similar to that of the mid-20th century.*

IN 1964, having taken a year off from college, I got a job in the General Motors pavilion at the New York World's Fair in Flushing Meadows. Sporting a red blazer emblazoned with the Cadillac crest, I escorted VIP guests on the Futurama II ride, a successor to the original 1939 Fair's Futurama, billed as "a journey for everyone today into the everywhere of tomorrow." To a sonorous narration, moving chairs carried visitors past dioramas depicting an imminent automotive Zion. In one, a vast road-building machine used lasers to carve highways through the Amazon Jungle, "bringing to the innermost depths of the tropic world the goods and materials of progress and prosperity."

In another, robots worked the oil deposits of the continental shelf. On the continent, cities multiplied, teeming with electronically paced traffic and bristling with soaring skyscrapers and underground parking garages. When humans were visible, it was as tiny stick figures without discernible clothing or features.

Elsewhere on the fairgrounds, other corporations presented their own tomorrows. GE's Progressland featured the Disney-designed Carousel of Progress, which traced the history of an audio-animatronic family from the 1890s to modern push-button living. In an egg-shaped pavilion designed by Eero Saarinen, IBM demonstrated machine translation and handwriting recognition. Bell debuted its picture phone — and a troupe at the Du Pont Pavilion celebrated the miracles of modern chemistry (Tedlar! Fabilite! Corfam!).

For the visitor, it was one-stop shopping for all the corporate versions of the technological sublime — the expression of a uniquely American conviction that technology would, as Leo Marx put it, "create the new garden of the world." In earlier American apostrophes to the machine, awe was tinged with the terror that was essential to the Burkean sublime. Walt Whitman sounded that note in his paean to the "fierce-throated beauty" and "lawless music" of the steam locomotive in *Leaves of Grass* — as did Arturo Giovannitti in his "The Day of War" (1916), with its modernist description of the skyscraper as "[challenging] the skies, terrible like a brandished sword."

But there was none of that here. Visitors to the Futurama saw nothing unsettling as they looked down at the miniaturized road-builders and skyscrapers. "Technology can point the way to a future of limitless promise,"

the narration concluded. But what it presented was hardly limitless, just a familiar, domesticated, consumerist tomorrow — one that was "awesome" only in the degraded modern sense of the word. Like the Jetsons (then in their second season), we'd keep on doing what we always had, only more efficiently.

¤

This realization figures prominently in Nicholas Carr's collection *Utopia Is Creepy*. The selections are mostly posts from Carr's tech blog *Rough Type*, written between 2005 and 2015. The book is fleshed out with a nosegay of tweet-sized aphorisms and a few longer essays that contain some of the book's best writing. The result is what has been called a blook (a term that inspired one self-publishing platform to launch a short-lived "Blooker Prize"). The genre has its limits. A blog entry is inevitably less compelling when it appears bare and linkless on the printed page, years after its posting — particularly when the topic is technology. If, as people say, one internet year corresponds to seven calendar years, then the earliest selections in this collection go back to the digital equivalent of the Truman presidency. It's hard to work up any interest in Carr's thoughts about Steve Jobs's presentation of the first iPhone or the controversies over the commercialization of *Second Life*.

Still, a long view of the period can be useful. If nothing else, it gives us an idea of what a sprawling landscape the label "tech" encompasses. Over the years, Carr's posts have touched on, among other things, social media, search engines, open source, Wikipedia, high-frequency trading, wearables, Big Data, self-tracking devices, smartphones, AI, video games, music streaming, and holography. Yet, as Carr's compilation makes clear, even as technologies replace one another and vogues come and go, the rhetoric of progress simply readapts. As at the Fair, each new technological wave throws off its own utopian vision. The posts give us glimpses of an internet Eden, a Big Data Eden, a cyborg Eden, a biosensor Eden, an automation Eden. Carr has taken on some of these scenarios at length in earlier books like *The Shallows* and *The Glass Cage*, which established his reputation as a thoughtful critic. But it's instructive to see him summarily dispatching them one after another here. He's hardly a Luddite, at least in the loose modern sense of the term: someone with a naïve and unreasoning aversion to machines. Rather, he can be eloquent and engaging when discussing the human costs of technology and automation, or the self-serving delusions of enthusiasts.

¤

Technological utopianism is always self-aggrandizing. "We stand at the high peak between ages!" the poet Filippo Tommaso Marinetti wrote in his "Manifeste du Futurisme" in 1909, predicting, among other things, that the Futurist cinema would spell the end of drama and the book. Every other modern era has seen itself in exactly the same way, poised at the brink of an epochal transformation wrought by its newly dominant technology, which, as Carr notes, is always seen as "a benevolent, self-healing, autonomous force […] on the path to the human race's eventual emancipation."

At the moment Carr started his blog, the agent of millenarian change was the internet — in particular, what enthusiasts were touting as "Web 2.0," with its promise of universal collaboration, connectedness, and participation. User-created content like Wikis and blogs would displace the old media, and participants would ultimately congeal into a collective intelligence capable of acting on a global scale.

Carr's blog first came to wide attention on the strength of his critique of an influential article called "We Are the Web," by *Wired*'s "Senior Maverick" Kevin Kelly. Kelly wrote that the accumulation of content on the web — from music, videos, and news, to sports scores, guides, and maps — was providing a view of the world that was "spookily godlike." By 2015, he predicted, the web would have evolved into "a megacomputer that encompasses the Internet […] and the billions of human minds entangled in this global network." With chiliastic zeal, he announced, "There is only one time in the history of each planet when its inhabitants first

wire up its innumerable parts to make one large Machine […] You and I are alive at this moment." Future generations, he said, will "look back on those pivotal eras and wonder what it would have been like to be alive then." Or, as Wordsworth might have put it, "Bliss was it in that dawn to be online."

In a post called "The Amorality of Web 2.0," Carr taxed Kelly with using a "language of rapture" that made objectivity impossible: "All the things that Web 2.0 represents — participation, collectivism, virtual communities, amateurism — become unarguably good things, things to be nurtured and applauded, emblems of progress toward a more enlightened state." On the contrary, he countered, those features are invariably mixed blessings. As a manifestation of the age of participation, Wikipedia is certainly useful, but it's also slipshod, factually unreliable, and appallingly written. "It seems fair to ask," he said, "when the intelligence in 'collective intelligence' will begin to manifest itself."

Similarly with blogs: Kelly described them as part of "a vast and growing gift economy, a visible underground of valuable creations" that turns consumers into producers. Carr, himself a blogger, pointed to the limits of the blogosphere: "its superficiality, its emphasis on opinion over reporting, its echolalia, its tendency to reinforce rather than challenge ideological polarization and extremism." In short, "Web 2.0, like Web 1.0, is amoral. It's a set of technologies — a machine, not a Machine — that alters the forms and economics of production and consumption."

Carr's post was widely discussed and contested. In retrospect, it seems largely unexceptionable. True, the promises of participation and amateurism have been realized. We're awash in user-created content, including not only Wikipedia but also Facebook and Twitter, Yelp and Reddit, Pinterest and Flickr, and all the rest. But when data scientists run their analytics on those media, they find the interactions they enable to be extraordinarily fragmented and polarized. And what glimpses we get of the whole are apt to curl our hair. (As Clay Shirky observes: "The internet means we can now see what other people really think. This has been a huge huge disappointment.") Worse still, in public discourse, the voices of amateurs, while plentiful, are increasingly drowned out by more-or-less corporate ones. Individual blogs like Carr's are still around, but no one speaks of the Blogosphere anymore; rather, blogs exist cheek by jowl with a variety of new media like *Mashable*, *Huffington Post*, and the content octopus *BuzzFeed*, as well as the online operations of old ones. As Carr points out in another post, branded content has driven the peculiar and idiosyncratic to the bottoms of search pages and the margins of our attention — "the long tail has taken on the look of a vestigial organ."

The persistent question is not "when the intelligence in 'collective intelligence'" will begin to manifest itself, as Carr asks, but when are we going to get to the "collective" part? You can look in vain for a perch that provides a "godlike view" of the whole, or for an emergent "community of collaborative interaction." We thought we were building the New Jerusalem but we wound up with something more like the current one.

¤

Anyway, the conversation has moved on, as it always does. Looking back over the history of technological enthusiasms in his *American Technological Sublime*, the historian David Nye notes that, in each generation, "the radically new disappears into ordinary experience." By now, the internet is ubiquitous, and for just that reason no longer a Thing. There are between 50 and 100 processors in a modern luxury car, about as many as there are electric motors (think power steering, seats, wipers, windows, mirrors, CD players, fans, etc.). But you wouldn't describe the automobile as an application of either technology.

So the futurists have to keep moving the horizon. One feature that makes this era truly different is the number of labels that we've assigned to it. Carr himself lists "the digital age, the information age, the internet age, the computer age, the connected age, the Google age, the emoji age, the cloud age, the smartphone age, the data

age, the Facebook age, the robot age"; he could have added the gamification age, the social age, the wearable age, and plenty of others. Whatever you call it, he notes, this age is tailored to the talents of the brand manager.

In his more recent posts, Carr is reacting to these varying visions of a new millennium, where the internet is taken for granted and the transformative forces are innovations like wearables, biosensors, and data analytics. The 2011 post from which he draws his title, "Utopia is creepy," was inspired by a Microsoft "envisionment scenario." Direct digital descendants of the World's Fair pavilions, these are the videos that companies produce to depict a future in which their products have become ubiquitous and essential, similar to the worlds pervaded by self-driving cars or synthetics described above. The Microsoft video portrays "a not-too-distant future populated by exceedingly well-groomed people who spend their hyperproductive days going from one computer display to another." A black-clad businesswoman walks through an airport, touches her computerized eyeglasses, and a digitized voice lights up to define a personal "pick up" zone:

> As soon as she settles into the backseat the car's windows turn into computer monitors, displaying her upcoming schedule […] [h]er phone, meanwhile, transmits her estimated time of arrival to a hotel bellhop, who tracks her approach through a screen the size of a business card.

One thing that makes these scenarios disquieting, Carr suggests, is the robotic affectlessness of the humans — who bring to mind the "uncanny valley" that unsettles us when we watch their digital replicas. These figures are the direct descendants of those audio-animatronic families that Disney designed for the 1964 World's Fair. As technologies become the protagonists of the drama, people become props. The machines do the work — observing us, anticipating our needs or desires, and acting on what they take to be our behalf.

It's that sense of ubiquitous presence that has made "creepy" our reflexive aesthetic reaction to the intrusiveness of new technologies — there is already a whole body of scholarly literature on the subject, with journal articles titled "On the Nature of Creepiness" and "Leakiness and Creepiness in App Space," etc. Creepy is a more elusive notion than scary. Scary things set our imaginations racing with dire thoughts of cyberstalkers, identity thieves, or government surveillance. With creepy things, our imagination doesn't really know where to start — there is only the unease that comes from sensing that we are the object of someone or something's unbidden gaze.

That creepy note is endemic to the enthusiasms of the Quantified Self (think of wearables like Fitbits and personal trackers). Like many new technologies, writes Carr, these were originally envisioned as liberating but have wound up as instruments of social control. He mentions the Hitachi Business Microscope, a sensor worn by employees on a lanyard, which monitors their movements, their interactions, and how often they speak up at meetings, all in the interest of "contributing to organization productivity by increasing the happiness of a group," as Hitachi puts it. It's symptomatic of what Carr calls a New Taylorism; these tools extend and amplify the reach of employee measurement from Frederick Winslow Taylor's time-and-motion studies of factory workers to the behavior of white-collar ones, who are already under surveillance by software that registers their every keyboard click. In the modern case, though, the supervisor is a hovering but unseen presence, unobservable on the jobsite.

¤

What's most striking about these pictures of the sensor-saturated world isn't just their creepiness, but how trivial and pedestrian they can be. The chief of Google Android touts interconnected technology that can "assist people in a meaningful way," and then offers as an example automatically changing the music in your car to an age-appropriate selection when you pick up your kids. Microsoft's prototype "Nudge Bra" monitors heart rate, respiration, and body movements to detect stress and, via a smart phone app, triggers "just-in-time

interventions to support behavior modification for emotional eating." (A similar application for men was judged unfeasible since their underwear was too far from the heart — "That has always been the problem," Carr deadpans.) They're symptomatic of Silicon Valley's reigning assumption, writes Carr, that anything that can be automated should be automated. But automatic music programming and diet encouragement — really, is that all?

Others extend these technologies to scenarios in which everything is centralized, rationalized, and Taylorized. In the futuristic reveries of PayPal's co-founder Max Levchin, the ubiquity of networked sensors — in the world and in our bodies — will make it possible to eliminate the chronic inefficiencies of "analog resources" like houses, cars, and humans. That's exactly what companies from Walmart to Uber are doing right now, but this only scratches the surface. Why not introduce "dynamically-priced queues for confession-taking priests, and therapists"? Why not have a car seat equipped with sensors that can notify your insurance company to increase the day's premium when you put your toddler in it, and then reduce it when it turns out you've only driven two miles to the park. Levchin even contemplates maximizing the power of the human mind. Imagine dynamic pricing for brain cycles. With brain plug firmware installed, why not rent out your spare cycles while you sleep to solve problems like factoring products of large primes?

But, as Carr notes, if your insurance company can adjust your premium according to who's in your car, it can also adjust it according to how many slices of pizza you eat. This is the nightmare world of Big Data, he says, where "puritanism and fascism meet and exchange fist bumps." But Levchin offers it as a utopian vision, where technology can improve people's lives by delivering "amazing opportunities they wouldn't have in today's world." True, it introduces new risks — Levchin mentions bias and threats to privacy, and it hardly stops there — "but as a species," he adds, "we simply *must* take these risks, to continue advancing, to use all available resources to their maximum." As Carr notes, it's that conspicuous willingness to break a few eggs that enables tech visionaries like Levchin to cast themselves as the protagonists of a heroic narrative. Yet when it comes to the crunch, all that centralized control doesn't promise to make the future more exciting, just more efficient, like the electronically paced traffic on the Futurama freeways or the push-button kitchens in GE's Progressland. To paraphrase what Karel Čapek said about intellectuals, has there ever been anything so awful and nonsensical that some technologist wouldn't want to save the world with it?

¤

Geoff Nunberg teaches at the UC Berkeley School of Information. He writes on language and technology and is often heard on NPR as Fresh Air*'s Language Guy. His latest book is* Ascent of the A-Word: Assholism, The First Sixty Years *(2012).*

10

Creators of the World Unite

by McKenzie Wark (November 25, 2015)

Reviewing Cory Doctorow's Information Doesn't Want to be Free*, McKenzie Wark draws on Doctorow's history of struggles over information control to provide the seeds of a "class-based theory of creative production" for the 21st century.*

A CANADIAN VERSION OF a "red diaper baby," Cory Doctorow is exceptionally well-placed to figure out what is and isn't in the interests of artists and creators in today's media landscape. Born in 1971 to Trotskyist militants with advanced degrees in education, Doctorow dropped out of college to become a computer programmer, founded a software company, and then went to work for the Electronic Frontier Foundation. In 2001 he became an editor at boingboing.net, one of the first successful weblogs, focusing on the intersection of technology and the counterculture. (A typical headline might read "3-D Printer Creates DRM-free DIY Steampunk Unicorn Tattoos.") He is also the author of a number of science fiction novels, for both adult and young adult readers.

Doctorow's creative credentials are in order, then, and in his new nonfiction book *Information Doesn't Want to Be Free* he has produced an essential primer for artists seeking to navigate the shark-infested waters of today's media. *Information Doesn't Want to Be Free* is more than a how-to book, however, or a *Creators' Rights for Dummies*: it provides the elements for a class-based theory of creative production in the 21st century.

A number of writers — including Douglas Rushkoff and Astra Taylor — have recently published useful books on similar issues, but Doctorow comes closest to understanding the interest of creators as a *class* interest. In my view, the class perspective runs like this: Creators are different from other workers in ways that make it difficult to analyze their behavior and interests in the class-based terms devised to make sense of 19th-century industrial labor. For one thing, musicians, writers, sculptors, and video makers all view themselves as doing very different things, and marketing the products of their labor within very different industries. Their work is also difficult to standardize: where traditional laborers make the same thing, over and over, creators, by definition, make different things: the whole point of being a creator, and the value of what the creator makes, is that each product comes out at least a little bit different. While creators, like Marx's workers, rarely own the means of production, they also don't usually sell their labor power itself. Instead, they sell the rights to reproduce their work and make money from it.

Things become even more confusing when you consider that a lot of what creators make these days takes the form of information, as opposed to physical objects. The product of their labor is a file you can copy,

regardless of whether the file is text or sound or images or moving images. Audiences, like creators, usually don't own the means of making and distributing creative work; they need platforms and software provided by tech companies and service providers to access the creator's work at all. The overwhelming fact about life in this overdeveloped world of ours is that we don't make our own culture for each other. There's a whole host of culture and communication industries that stand between the creator and the audience.

The relationship of creator to audience is mediated by what Doctorow calls "investors" and "intermediaries," typified by Hollywood (the culture industry) and Silicon Valley (the internet-based communication industry). The investor owns a stock of creative works — images, texts, recordings — most often by controlling the copyright. The intermediary controls the flow of information: they own the means of getting copies of those books or songs to audiences. Sometimes the investor and intermediary are parts of the same multinational company, and often they are not.

Investors, when they can, will try to jack up the price of reproducing creative work. Intermediaries, when they can, will take advantage of creative work they can get for free. They will also try to monopolize the channel, inflate the price to the audience, and low-ball what they pay to the investor. Since creators have the least bargaining power of everyone involved in these transactions, they usually get the rough end of the pineapple.

Information Doesn't Want to Be Free provides a helpful thumbnail sketch of the history of these struggles over who benefits from creative production. Here's how it came about, in miniature: The piano roll business got started back in the 1880s by pirating sheet music. Sheet music investors objected, and a form of compulsory license arose, obliging the piano rollers to pay a set fee. The piano rollers went legit, and became the recorded music business. Then, in the 1920s, radio came along and began pirating recorded music. The music publishing investors objected, and a compulsory license arose, obliging radio stations to pay a set fee to play each song. The radio stations went legit, and became the broadcasting industry.

The recurrent pattern is one of new intermediaries getting started through the exploitation of creative works they got for free, followed by investors fighting back, followed by government stepping in to regulate the relation between investor and intermediary. A new intermediary gets started by offering an audience a great new way to access stuff, but part of what makes it work is that the intermediary doesn't have to pay for the stuff. The intermediary cuts out the investor; the investor fights back. Then the government intervenes by imposing some sort of compulsory licensing regime in which intermediaries can stick whatever they want in their tube but have to pay a fixed fee of some sort to do so. This is supposed to be distributed among the investors, who are supposed to pay some of that back to the creators. As a side effect of these state-refereed wrestling matches between intermediaries and investors, creators sometimes manage to make a little money, too.

¤

Compulsory licenses have their flaws, as Doctorow acknowledges. A compulsory license always involves a trade-off, from the creator's point of view. It's a system where anyone can copy your work so long as they pay the license. This is how bands get permission to cover songs written by other bands, for instance: they pay a fee to a collection agency like ASCAP and the song becomes effectively (temporarily) theirs. The original songwriter might hate the cover version, but so long as the appropriate fee is paid, there's nothing she can do about it. So there's a sacrifice of a so-called "moral right" of the creator here, in the interest of their financial remuneration. Most creators are happy with that trade-off.

Doctorow makes a good case for extending the compulsory license model further to cover the current wave of technological change. You want to remix a recording of a pop song? You want to stream your favorite TV show over a network to your laptop so you can watch it while you're away from home? Sure, why not? Should

be easy. It is not hard to imagine a world in which creators get paid for their work, and yet audiences have a lot of flexibility about how they receive their cultural stuff and what they can do with it.

Investors and intermediaries hate this idea. The former are not really in the business of investing in the best culture. The latter are not really in the business of building the best networks. Both are in the business of making money by controlling, respectively, stocks or flows of information. So naturally they want to extract a rent — a super-profit — from that control wherever they can. Investors in information stocks have it in their heads that they can charge for every instance in which their stock is used, ever, for all time. Intermediaries who control information flows believe that they can lock both investors and audiences into their proprietary pipelines, thereby gaining an advantage over both as well as their rival intermediaries.

Both investors and intermediaries have spent a lot of time, money, and energy trying to recruit creators and audiences to their respective points of view. The investors tell creators that the intermediaries are evil, making your stuff available for free and stealing your money. (YouTube, Spotify, Amazon: these are some of the bogeymen in this narrative.) The intermediaries tell creators that the investors are evil, acting like monopolists and manipulating prices. (Movie studios, record companies, the Big Five publishers: another rogue's gallery.) The investors frighten audiences into thinking the intermediaries are conduits for child porn or terrorist propaganda. The intermediaries warn audiences that what the investors really want is some sort of Chinese-style police state surveillance of your streaming and browsing habits, all to stop teenagers uploading videos of themselves dancing in their bedrooms to a pop song. All of which, it transpires, is hardly exaggeration. Both investors and intermediaries really do pursue such extreme tactics, and neither truly have creators' interests at heart. A plague on both their houses.

¤

There's an ever-present tension between how the culture industry works and how actual culture works. The culture industry wants to make all creative output into a commodity. Everything is to be reduced to its exchange value. But culture has never worked like this, at least not entirely. It is also, as commentators from Marcel Mauss to Lewis Hyde have noted, a world of gift-giving. I play a record for a friend because I think they might like it, and they do, so I lend it to them. I may never get a return on this "investment." Maybe my friend will forget to give the record back. And later invite me to the movies — his treat.

Back in the days of mechanical reproduction, a lot of the sharing aspect of culture was invisible. A lot of what went on in culture took place outside the commodity economy and, strictly speaking, outside the law, taking place beneath the radar of investors and intermediaries. Record labels weren't paying private investigators to follow their customers around and make sure they didn't lend albums to each other.

Investors don't recognize that gift-giving is part of what culture is all about. Certainly their lawyers don't. Lawyers get paid to read or to listen or to watch. Every minute is billable. The idea that the rest of us share cultural information as a gift to each other, for free, out of love, just doesn't make sense to them. Thus, the copyright lawyer's obsession with locking down any and every possible use of a cultural artifact. Everything that is not a sale must be a lost sale: "piracy."

In the late 20th and early 21st centuries, investors took to attacking the gift-giving side of culture technologically as well as legally. They thought they were being hard-nosed by protecting their interests, but this turn to technology, as Doctorow shows, was how the intermediaries managed to play them for suckers. The con was — and is — DRM, Digital Rights Management, or what Doctorow calls the "digital lock." It's an encryption code that prevents the content of a file from playing on machines that lack the right digital cryptographic key.

The digital lock, Doctorow argues, was a decisive win for intermediaries. Investors wanted to lock the audience *out* of the technological means of copying, but ended up locking themselves *in* to the proprietary systems of the intermediaries. So the song or the book or the movie is now stuck as a useless file, unless you have access to the special product from Apple or Amazon on which it can be accessed. The intermediaries now have everyone right where they want them: in locked channels that they alone control.

It gets worse: because the digital locks don't work very well and are easily hacked, they have to be continually updated. That means the intermediaries have to gain remote control over your computers so that they can periodically change the locks. You are not allowed to interfere with this process, or even know how it actually works.

Clearly, this is bad for the audience. But the investors have their own ideas about what is in their interests, which often conflict with those of everyone else. Investors are pretty obsessive protectors of their monopolies over information stocks, and they want intermediaries to be legally liable for every single unauthorized copy that ever gets distributed over their networks. That way investors can get rich suing intermediaries, or drive out of existence those they do not control.

Now, if I make an illegal physical copy of something and send it via FedEx to everyone in my address book, FedEx is not in any legal trouble. It is simply the common carrier. So long as FedEx agrees not to play favorites and will ship anything for anyone, it isn't responsible when its customers ship something that breaks a copyright. That's a matter between the investor who owns the copyright and the infringing party — me! — not FedEx.

It would appear to make perfect sense to extend this sort of logic to the various forms of internet intermediaries. So long as your internet service provider or your social media site is abiding by certain fair-play rules, why should it not have the same sort of common carrier status? There is already a simple process for dealing with copyright violations, after all: the owner of the copyright sends a notice to the party in violation, and so long as the infringing information comes down, there's no additional liability on the part of the intermediary.

As Doctorow notes, this is a system that is open to abuse. If someone puts something about your company on the internet that you don't like, you can just claim a copyright violation. But the take-down notice system, whatever its flaws, is at least a bit better than what investors really want. They want to have access not only to public but also to all private communications to look for copyright violations. And they want intermediaries to be held responsible for any violations that occur on their networks.

It's easy to see why investors want the intermediaries to be liable. Nobody wants to be the bad guy, and it really doesn't look good when some giant media conglomerate takes a poor teenager to court for downloading a few of its pop songs. Besides, the kid probably doesn't have any money anyway — which rather punctures the lawyers' argument about lost sales. Investors would rather sue intermediaries, which tend to be large corporations with deep pockets.

The power to look into private communications is something else again. As Doctorow points out, if he wants to use YouTube to send a video of his child in the bathtub to relatives, how is that any of the investors' business? So what if some Taylor Swift song was playing in the background? In short, the investors want the whole communication system to be redesigned around their particular problems.

¤

It is tempting, when considering all of these Machiavellian stratagems, to take a side and support the investors

against the intermediaries or vice versa. But — and this is crucial to bear in mind — their interests are not our interests. As Doctorow sagely notes:

> The future of the Internet should not be a fight over whether Google (or Apple or Microsoft) gets to be in charge or whether Hollywood gets to be in charge. Left to their own devices, Big Tech and Big Content are perfectly capable of coming up with a position that keeps both "sides" happy at the expense of everyone else.

So what's a creator to do? Well, the first thing might be to stop believing that either the investors or the intermediaries have your back. Investors want to treat what you make as private property to be annexed for their vast fiefdoms, from which they think they are entitled to a rent in perpetuity. Intermediaries want to charge everyone for access to your work as it passes through their pipelines, and gather a lot of additional information about them in the process. Neither care in the least about your need to make a living or the contextual integrity of your work. Therefore, Doctorow thinks that canny creators need to think tactically about when their interests can be aligned with those of investors or those of intermediaries, or neither.

An example of an issue on which creators need to pick their allies carefully is net neutrality. Internet intermediaries are creatures of tax giveaways and regulation; without the largesse of the state, they couldn't exist. And yet they want to claim they should be "free" to charge premiums simply for delivering your information in a timely fashion. If internet intermediaries are to be treated as common carriers, exempt from copyright liability, given right-of-way to lay their cables and so forth, then the least that can be asked of them is that they not discriminate between customers.

To give another example of the collective calculation of interests needed by creators as a class: In Doctorow's opinion, it's better tactics for creators to side with intermediaries who let people access stuff for free than those who advocate for digital locks or for some surveillance-based copyright enforcement regime. It's a question of priorities: the artist has to first think about becoming known before she thinks about getting paid. Having information circulate freely at least helps clear the first of those obstacles, even as it does little to address the second.

This is not the first time creators have had to revise their strategies to deal with technological change. Doctorow describes how musicians of the 1930s who did not particularly want to record but rather focused on live performance found themselves outflanked by those who invested time and effort in record-making. They then got to tour even more because broadcasting made their recorded sound popular. Now things are rapidly going the other way. Musicians who would prefer to just record and don't (or can't) play live in front of audiences find it harder to gain an audience or make a living without touring. As always, technology's pendulum swings favor some and devastate others.

Doctorow makes a persuasive case that it is generally in the creator's interest for art to circulate online for free, as a kind of gift. But if you want to have a class of artists able to support themselves on the proceeds of their creative activity, you still have to come up with a plan to commodify part of the art. You can sell tickets for in-person appearances, like the touring musician does. You can sell swag, like T-shirts or coffee mugs. There's advertising, which is what has supported boingboing.net since 2003. You can crowdsource your livelihood, offering special premiums to backers on Kickstarter or appealing to the goodness of your fans' hearts via Patreon. You can do work on commission. There are other options, which Doctorow leaves out, like making one-off objects for wealthy patrons, otherwise known as the art market. Or you can get a day job explaining your art in academia — if you can find one.

Frankly, none of these options are great, but that isn't really the fault of the internet. The options have never

been all that great for artists. If the public isn't purchasing your masterworks, it may not be because they can download them for free: it's more likely they just don't care for them, or even know they exist. No tactics or technology can counteract the power of public indifference.

¤

Doctorow is no techno-utopian, but he does have a political or ethical lodestar that orients his work: "Anything that minimizes the drag on our collective efforts should be celebrated." There are tools out there now, he tells us, that can support our common life. As he puts it: "Information doesn't want to be free, people do."

If you are inclined to think, as I am, that we are all cyborgs — perplexing mixes of flesh and tech — it might not be possible to divide "information" and "people" so cleanly. Once information started to be produced as something relatively autonomous from the material substrate that sustained it, it really does appear to have some agency. It seems to "want" to be shared like a gift. As I put it in a different riff on the same meme: information wants to be free but is everywhere in chains. It keeps getting stuffed back into the narrow confines of rather old-fashioned models of absolute private property.

The fundamental problem, I think, has to do with information itself. It is always somewhere in between complete chaos and complete order. It is neither noise nor a constant tone. It's always a pattern with lots of redundancy. Take this article: I made up none of these letters, none of these words. These are old patterns I am repeating; I just changed the order of the words around a bit. Yet from the point of view of law and the market, I alone am its "creator."

Hence it's always rather artificial to call any particular piece of information someone's "property." It's always just a bit of a larger pattern of information that has been collectively made. It really does want to be free — it doesn't know how not to be — but, alas, there is a class of people who interpose themselves between us and our information, by controlling stocks or flows of it. They seem to think everything about its potential for common life has to be sacrificed to their desire to keep collecting the rent.

Doctorow is one of the finest organic intellectuals of our time, and *Information Wants to Be Free* is a valuable and provocative book that should be read by creators and audiences alike. (Investors and intermediaries might want to stay away.) As I have tried to show, it also includes the seeds of a class analysis that is sorely needed to counter both the techno-phobic and techno-libertarian prejudices that are so prominent in today's discourse about creative industries and the internet.

I have etched the theoretical outlines of the class conflict of creators versus investors and intermediaries rather more sharply than Doctorow does, while he more profoundly details these issues on the technical and political level. His book provides a handy reference on a wide range of such matters, and we need that, too. But I think it has the added benefit of offering food for thought about the long-term historical struggle over information: a struggle that belongs to all of us, whether we accept it or not.

¤

McKenzie Wark is the author of A Hacker Manifesto *(Harvard, 2004),* Molecular Red: Theory for the Anthropocene *(Verso, 2015), and various other things. He teaches in the Liberal Studies MA at the New School for Social Research in New York City.*

11

Automating the Professions: Utopian Pipe Dream or Dystopian Nightmare?

By Frank Pasquale (March 15, 2016)

Frank Pasquale argues that, in their book The Future of the Professions: How Technology Will Transform the Work of Human Experts*, Richard and Daniel Susskind overstate the impact of automation on healthcare, the law, and other professions.*

RICHARD SUSSKIND has been discussing "the end of lawyers" for years. He's at it again, but this time with even more sweeping claims. In a recent book entitled *The Future of the Professions*, co-authored with his son, Daniel, he argues that nearly all professions are on a path to near-complete automation. Lawyers may be relieved by this new iteration of his argument — if everyone's profession is doomed to redundancy, then law can't be a particularly bad career choice after all. To paraphrase Monty Python: few expect the singularity.

The Susskinds don't just foresee the robotization of accountants, actuaries, and architects. They anticipate the emergence of bot preachers, ministering to avatars in *Second Life*'s virtual cathedrals. They ride the "massive open online course" bandwagon, however stalled it may be at present. They argue that struggling newspapers may well use Narrative Science algorithms (now deployed to churn out earnings reports and sports news) to write up politics, business, and style stories. And they quote a study by Carl Benedikt Frey and Michael Osborne arguing that only one percent of tax preparation work is safe from computerization. The future of the professions is thus, for them, at the mercy of artificial intelligence–driven (AI) startups.

How realistic are such scenarios? The book argues that even the most complex, professionalized work is on the road to computerization. The logic is straightforward: machines will watch what professionals do, record their "outputs," and reproduce them. But, in fact, this logic is far too simple. The Susskinds cherry-pick the Frey and Osborne study, highlighting the parts they like while studiously ignoring its more optimistic conclusions about surgeons, nutritionists, clergy, teachers, and many other professionals. Indeed, the Frey and Osborne study is far less negative about the viability of the professions than the Susskinds' book, assigning a "low" probability of computerization to myriad professional positions:

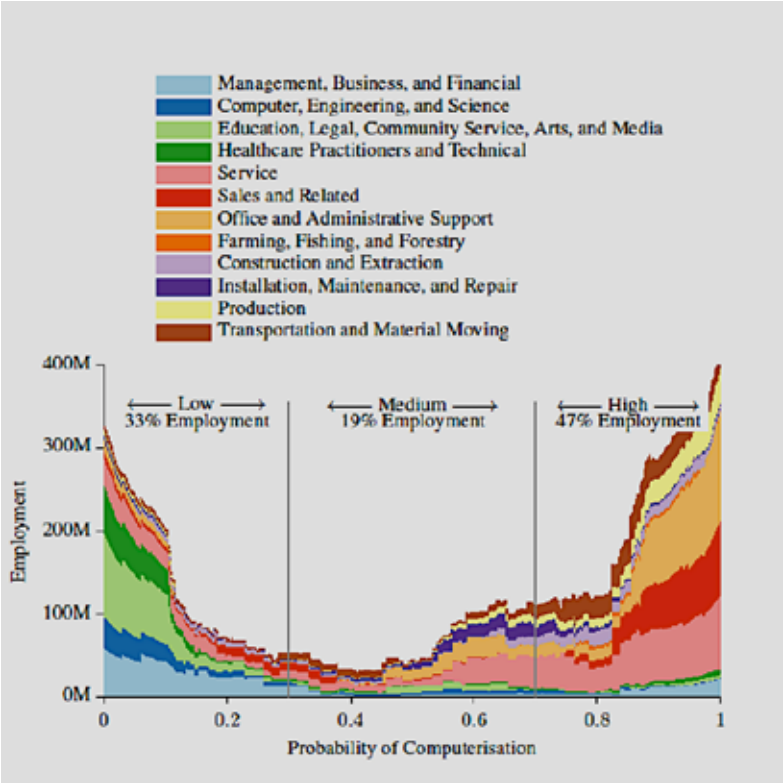

It is hard to disprove Frey and Osborne's conclusions without exaggerating the quality of existing data recording, the power of current or near-future algorithms to integrate data streams, and the ability of engineers with little or no domain expertise in a given profession to code software replacing its human experts. Sadly, the Susskinds fall into these traps even where one would presume their expertise would be greatest — in law, where Richard Susskind has had some influence.

Legal Automation: Symbiosis, not Substitution

The Susskinds welcome the accelerating automation of litigation and transactions. They claim that the form provider "LegalZoom operate[s] reliably and to a high level of user satisfaction," and they extrapolate from its success (and that of firms like it) to the rest of the profession. Unfortunately, they fail to credibly estimate how much of extant legal work is as automatable as, say, drafting a very basic will. Researchers Dana Remus of UNC and Frank Levy of MIT, who have done that work, project far milder effects on the profession. Many of the new software-driven "legal services" take on low-value tasks that were either rarely performed by lawyers or rarely generated much revenue for them (like uncomplicated contract drafting). Other software services might very well generate, rather than destroy, complex work for lawyers. For example, novel procedural moves accelerated by software may empower a whole new generation of plaintiffs' lawyers.

Experts differ on the real likelihood of pervasive legal automation, and its effects. Frey and Osborne put the risk to lawyers at under four percent, and project that paralegals are in much more danger. But empirical research by economist James Bessen has complicated even that assumption: "Since the late 1990s, electronic document discovery software for legal proceedings has grown into a billion dollar business doing work done by paralegals, but the number of paralegals has grown robustly." Like MIT's David Autor, he observes that automation can create more jobs than it eliminates. Considering the role of text-mining software in law firms, the management consulting firm McKinsey says automation amplifies "the value of expertise by increasing an individual's work capacity and freeing the employee to focus on work of higher value." Researchers Michael

Simkovic and Frank McIntyre reach a similar conclusion. All of this research taken together suggests that the Susskinds commit the classic "lump of labor" fallacy, assuming that there is some fixed amount of legal work to be done, and thus any cession of that lump to machines results in less work for extant professionals.

The Susskinds are not unaware of such research, conceding that "there is no fixed 'lump' of labour." They nevertheless insist that, even if new tasks arise, "machines might be better placed than people to undertake" them. But their subjunctive mode here can just as easily be reversed: people might be better placed than machines. We should thus expect corporations, governments, lawyers, and programmers to struggle over the scope and intensity of automation in coming decades. Context matters, and the Susskinds are all too prone to dwelling on harmless or uncontroversial automation while downplaying its more sinister overtones.

For example, consider red light cameras implemented without a right to appeal or dispute a ticket triggered by them. "The Department of Motor Vehicles could attach to any driver's license an "automatic debit permission" to routinize payment of fines." Such a system could completely standardize the law of red lights, and would of course render obsolete some portion of what current traffic attorneys do. But it would also make a mockery of due process and other core legal principles. Indeed, so-called "smart cities" might make many forms of law-breaking impossible — but could do so at great cost to protesters and the marginalized populations already besieged by discriminatory laws and policing. We would need to reform the law wholesale before blithely committing ourselves to its perfect enforcement.

Visions of future sociotechnical developments always reflect contestable political aspirations. It is easy to press for a decline in professionals' income or status by claiming that software could do their work. It's harder to explain why the many non-automatable aspects of their current practice should be eliminated or uncompensated. That's where the Susskinds falter, and where their book takes a rather nasty turn.

Repeatedly questioning the motives and competence of current professionals, the Susskinds insinuate that they resist change simply out of hidebound protectionism or technophobia. The *values* at stake in the human resolution of disputes — versus their automated dispatch to a computer system — are flattened into a utilitarian calculus of speed and efficiency. To counter critics of automation, they treat "cost-cutting" as an all-purpose trump card. The general public should be grateful for more legal automation, they argue, because human lawyers are too expensive.

This is cavalier, not to mention naive. I would worry about any person who decides to file a tort or contract case against a major corporation using an app. If the claim is frivolous, they could be sanctioned. If the claim is serious, it will probably be outmaneuvered by a (human) defense lawyer. And if corporations don't even need to deploy attorneys to deflect such interventions, but can even automate their own defense, then there's little reason to believe this will constitute some great triumph for justice. Our legal system exacerbates inequality because of uneven access to resources for advocacy, not lack of automation. Digital projects to "democratize the law" rarely include the more sophisticated predictive analytics the Susskinds trumpet; instead, they remain the exclusive preserve of wealthy corporations. The Susskinds give us little reason to believe that automation will impede — rather than accelerate — inequalities in legal resources.

Why is it so difficult to turn legal disputes over to a computerized simulation of justice? Situations involving conflicting rights, unusual fact patterns, and open-ended laws will remain excessively difficult to automate for the foreseeable future. As the former General Counsel of DuPont Canada, Seymour Trachimovsky, said of an earlier Richard Susskind broadside, predictions of pervasive legal automation exhibit "little appreciation of the huge volume of indecipherable legislation and regulation that is promulgated every day." *The Future of the Professions* argues that standardization of legal filings, briefs, and judgments is the way of the future, bringing the logic of interchangeable parts to the work of the mind. But, as Trachimovsky counters, "given

modern trends in legislative and regulatory drafting, in particular the use of 'creative ambiguity' […] demand for custom services will only increase."

Note, too, that much of this legal complexity is not simply a question of redundant, easily simplified legal verbiage. It reflects, instead, deep-seated social tensions that are unlikely to be resolved any time soon. For example, had America united behind the goal of universal healthcare in 2009, the Affordable Care Act could have been a sentence or two long: "1. Everyone is eligible for Medicare, and 2. Eligibility starts Jan 1, 2010." Instead, the Act went through myriad iterations as multiple stakeholders pressed their interests. It is now being interpreted in an epic series of rulemakings and adjudications on the federal and state level, which must address deep tensions between older models of healthcare regulation and the integrated delivery systems favored under PPACA. (For example: do we still forbid certain payments from hospitals to doctors as "kickbacks," or do we let the industry and physicians recharacterize them as "gainsharing incentives"?)

Not only is the complexity of the Act becoming ever more obvious, but so too is the fact that human values and judgment will be critical to its implementation. As just one example: Is it too much to ask a policyholder to travel 25 miles to find a gastroenterologist? 50 miles? Questions like this will determine the fate of "narrow networks" (a key feature of many exchange plans). If the Susskinds have an algorithm to "solve" such legal questions, I would love to see it — particularly given the extraordinary variation in geography, car ownership, and public transit access in the US.

The fate of artificial intelligence in law is far more open-ended than *The Future of the Professions* suggests. Only specific, situated, technical, and sociological analyses of particular areas of law are truly valuable here. The Susskinds hastily sketch the power of "sophisticated diagnostic expert systems, which tackle highly complex, multi-jurisdictional legal questions," but give us little sense of how their processing of tough legal questions is translated to clients, courts, or regulators. Presumably there are no direct brain downloads available yet, and each audience will demand more than a computational *ipse dixit*. Moreover, that demand for humanly intelligible explanation is important — no one should be denied access to benefits (or thrown into jail!) simply because a computer said yes or no.

Of course, technology has shaped, and will continue to influence, legal practice. But its effects can be checked or channeled by law itself. The Susskinds themselves finally arrive at this realization at the end of their book. Having extolled the virtues of technology at length, they hurriedly acknowledge the importance of law and policy in determining whether, say, automated life support systems should be able to terminate treatment for a patient who appears to be in a persistent vegetative state. They would like to believe that such ethical dilemmas are rare in order to retain the credibility of their central claim: that the vast majority of medical and legal scenarios can, and will, eventually be processed by software. But the medical field is rife with enduring legal and ethical dilemmas requiring difficult judgment calls and interpersonal communication.

Paging Dr. Google

Compared to the manufacturing, military, and finance sectors, the pace of automation has been sluggish in healthcare. Software has a difficult time mimicking the spontaneity, creativity, flexibility, and perceptiveness that are the hallmarks of good nurses, surgeons, psychiatrists, pediatricians, and so on. End-of-life care obviously requires that difficult decisions be made on the basis of imperfect information. The idea of mapping out every possibility in advance, and mechanizing the "termination of treatment," is bizarre for anyone with personal or professional experience in an intensive care unit, nursing home, or hospice. The stakes are so high, the concept of what makes life worth living so ineffable and variable, the decision-making so intricately shared among family and providers, that routinization seems ill-advised at best, and deeply offensive at worst.

Moreover, even elementary medical apps can fail patients. Earlier this year, the Federal Trade Commission settled lawsuits against firms who claimed their software could aid in the detection of skin cancer by evaluating photographs of the user's moles. The FTC argued that there was insufficient evidence to support such claims. The companies are now prohibited from making any "health or disease claims" about the impact of the apps on the health of users unless they provide "reliable scientific evidence" grounded in clinical tests. If algorithms designed merely to inform patients about their options aren't ready for prime time, why presume diagnostic robots are imminent?

The Susskinds have another easy answer at the ready: Moore's law. In other words, following Ray Kurzweil's reasoning in *The Age of Spiritual Machines*, they predict that computer processing speeds will keep dramatically improving; difficult problems will become trivial. But even if we concede that Kurzweil's controversial equation of human cognition applies to machine information-processing, the timeline of technological advance is necessarily fuzzy. Much of healthcare's big data is also "bad data"; it needs to be verified, standardized, and connected before software can do a decent job of analyzing it. That process will demand greater professional engagement with patient records, not less — both from researchers contemplating novel issues in data science, and attorneys and security professionals grappling with vexing privacy issues.

For every hour that software may save a radiologist by corroborating a difficult diagnosis, many more hours may be needed to integrate new data streams into a diagnosis, track the results of therapeutic interventions, and discover underlying genetic and environmental determinants of disease. New professionals will be needed, too — and not just as coders. Understanding how the body itself works (and not just correlations among computerized representations of it) will be critical.

Having considered the informational and diagnostic aspects of medicine, let's consider a third area: procedures. Yes, there are robotically assisted surgeries, albeit with a very limited scope of application at present. Automation optimists project that their early successes will quickly spread to other types of interventions. But MIT economist David Autor offers a general reality check about automation that applies with even more force here:

> Most automated systems lack flexibility — they are brittle. Modern automobile plants, for example, employ industrial robots to install windshields on new vehicles as they move through the assembly line. But aftermarket windshield replacement companies employ technicians, not robots, to install replacement windshields. Why not robots? Because removing a broken windshield, preparing the windshield frame to accept a replacement, and fitting a replacement into that frame demand far more real-time adaptability than any contemporary robot can approach.

Of course, futurists can probably imagine a robot in a self-driving car that can navigate itself to your car, drive it to a garage, and order other robots to replace the windshield. But even that scenario depends on a chain of contingencies. When the stakes are higher — for instance, replacing a kidney instead of a windshield — then even more back-up systems and planning will be necessary.

Even if hospitals deploy robot surgeons en masse, the ultimate "back-up system" would be a skilled human surgeon with some experience, flexibility, and creativity. Our aim should not be to replace such individuals, but to aid in their efficiency and effectiveness. The sequence and shape of automation in healthcare cannot simply be dictated from on high by engineers. Rather, domain experts need to be consulted, and they need to buy into a larger vision of progress in their field. Perhaps more of medicine should indeed be automated — but let's ensure that physicians themselves are lasting partners in that process. They should be helped, not replaced, by machines — both for the present (in order to override errant machines), and for the future (to develop new and better ones).

Expertise and Governance

The Susskinds address more than law and medicine. For them, education, architecture, journalism, management consulting, and even "divinity" all stand in automation's crosshairs. They briefly grapple with concerns about empathy and other "soft skills" in professions, but they argue that such traits of mind and character can either be translated into algorithms, or can be stripped out of traditional professional roles. As with their perfunctory responses to other objections, their treatment of empathy and human interaction is "almost entirely an exercise in rational deduction and accordingly devoid of historical and cultural considerations," as Howard Gardner explains in an insightful response to their work.

A persistent mistake undermines *The Future of the Professions*. The authors conflate the professional role with the delivery of expertise. Thus they fail to seriously address two issues at the core of professional identity. First, there is some degree of self-governance among professionals. They primarily work *with* clients or patients, for example, and not *for* bosses or shareholders. Second, the main reason they enjoy this autonomy is because they must handle intractable conflicts of values that repeatedly require thoughtful discretion and negotiation. In isolation, these factors damage the Susskinds' case; together, they prove fatal to it.

To start with the question of values: rarely, if ever, is a vocation simply a matter of conveying information. The duties of professionals do not end with an assessment of the relative likelihood of an intervention "working," where "working" is defined by a clear, quantifiable metric. That's most obvious in, say, elder care or education. A robot telling a shut-in elderly person, "Your friends and loved ones care for you" is not really a good substitute for visits. As for children in school, they require some guidance as to how to conduct themselves — as persons, not simply as absorbers and generators of messages. To think otherwise is to put society on a slippery Skinnerian slope to behaviorism.

There is no single measurement of success in the professions, and further complexities arise as soon as one lacks a single quantity or thing to be optimized. For example, attorneys, doctors, and teachers often face very difficult conflicts of values — between, say, zealous advocacy for a client, and acting as an officer of the court; or between extending a patient's life, and assuring quality of life in the time remaining; or between attending to disruptive students, or simply ordering them out of the classroom to ensure others can learn better. We mark the importance of these decisions by insisting that a human be directly responsible for them. Routinized or robotized approaches do not respect the dignity of the client, the patient, and the student.

Admittedly, both government and corporate bureaucrats will insist that costs can be cut by summarily resolving such issues. (Far be it from them to consider raising taxes on the wealthy to pay for the poor's access to human expertise.) The Susskinds' book is rife with technocratic language, and one can imagine it being invoked by big data and predictive analytics firms who promise to replace all manner of human judgments with "data-driven" scores. Unfortunately, both the data and code for such proprietary systems is rarely available to those outside the firms selling them. Black-boxed as trade secrets, such systems defy challenge, debate, and correction, which hardly constitutes a step toward the democratization of expertise promised by the Susskinds. Rather, such proprietary systems constitute a regression to intellectual feudalism, where small groups of plutocrats arrogate to themselves ever more control over expanding areas of human experience. *The Future of the Professions* pretends to anti-elitism, but its attitude and program would be entirely at home in the drawing rooms of Davos, Wall Street, and Silicon Valley.

In the Susskinds' hands, professionalism is too often misconstrued as the protectionism of guilds, dismissed as elitism, lambasted as a source of inertia. But sociologists like Eliot Freidson have observed that professionalism serves as a social and economic logic that tempers the worst excesses of free-market competition and cutthroat,

Amazon-style managerialism. Professionals have been granted some degree of autonomy because they are charged with protecting distinct, non-economic values that society has deemed desirable. Their labor, in turn, reflects, reproduces, and is enriched by those values. Knowledge, skill, and ethics are inextricably intertwined. We cannot simply make a machine to "get the job done," because frequently task definition is a critical part of the job itself.

Freidson lamented the frequent failure of professionals to "spell out the principles underlying the institutions that organize and support the way they do their work." The Susskinds' book deserves some credit for prompting professionals to better explain their roles and values. In the face of rampant and often reckless automation, the professions ought to reaffirm their own norms, highlight the importance of tacit skills and knowledge, and extend their status to other workers. The alternative is grim, and perhaps best described in Lawrence Joseph's recent poem "Visions of Labour":

> … the question in this Third
> Industrial Revolution is who owns and controls
> the data. That's what we're looking at, labour cheap,
> replaceable, self-replicating, marginal, contracted out
> into smaller and smaller units. Them? Hordes
> of them, of depleted economic, social value,
> who don't count, in any situation, in anyone's eyes,
> and won't count, ever, no matter what happens,
> the truth that, sooner than later, they will simply be
> eliminated. …

Joseph's chilling lines are a sparer, more honest account of the politico-economic developments celebrated by the Susskinds. If we are to maintain a democratic society rather than give ourselves over to the rise of the robots — or to those who, like the Susskinds, bid them to rise — then we must spread professionalism from areas like law and medicine to information retrieval, logistics, elder care, marketing, and many other fields. Imagine a labor movement built on solidarity among those who realize that, like bad bosses, poorly implemented automation can crush human potential and innovation. If they succeed in uniting, they might project a vision of labor far more concrete and realistic than the feudal futurism of techno-utopians. They might foster automation that complements extant skills and labor, rather than accelerates a cheaper, faster, and more catastrophically unequal version of the present. Perhaps the Susskinds' next book can explain how technology could advance the interests of all workers, rather than the narrow sliver of technologists, managers, and financiers likely to thrive in the future they commend.

When Pew polled hundreds of technology experts on the future of work, about half said that robots would render "masses of people […] effectively unemployable." *Guardian* columnist Jess Zimmerman had a quick retort: "Stop predicting the future. Leave that to the robot who's going to steal your job." Concrete assessments of the real progress of automation in the professions confirm the wisdom of more sober voices. Computers assisting, rather than replacing, professionals is the norm, and will continue to be so in any recognizably humane social order.

¤

Frank Pasquale is an expert on the law of big data, predictive analytics, artificial intelligence, and algorithms, and has served as a member of NSF-sponsored Council on Big Data, Ethics, and Society. His book The Black Box Society: The Secret Algorithms That Control Money and Information *(Harvard, 2015) developed a social theory of reputation, search, and finance.*

12

Bursting the Optimistic Technology Bubble

by Evan Selinger (July 31, 2015)

Evan Selinger critically examines Martin Ford's claim in Rise of the Robots: Technology and the Threat of a Jobless Future *that a techno-optimistic "bubble" is impeding smart discussion of the perils of information technology and automation.*

MARTIN FORD's *Rise of the Robots: Technology and the Threat of a Jobless Future* is a long-form exercise in dystopian scenario planning. Given how that genre reads, reviewers have naturally focused on this book's grim prognosis and radical solution for avoiding disaster. They've assessed the likelihood of innovation-driven capitalism collapsing under the weight of mass unemployment and decimated consumer purchasing power and confidence. And they've judged the cost and consequences of giving citizens a guaranteed minimum income as a safety net and as a risk-taking incentive.

Tempting as it is to offer my own predictions, I think other avenues need to be explored. In particular, there's still much to learn from Ford's technology-driven concern that "a great many people will do everything right […] and yet still fail to find a solid foothold in the new economy."

Ford has thought long and hard about automation. Back in *The Lights in the Tunnel: Automation, Accelerating Technology and the Economy of the Future* (2009), he cast aspersions on the so-called "Luddite Fallacy," a development economics concept often used to dismiss arguments about the end of work: "while technological progress will cause some workers to lose their jobs as a result of outdated skills, any concern that advancing technology will lead to widespread, increasing unemployment is, in fact, a fallacy." Ford moved the conversation about jobs in an interesting direction by questioning whether history is really caught in an ever-repeating loop.

Think of it this way: The market seems resilient because new jobs are created when gains in technological efficiency render older ones obsolete. As the much-repeated story goes, after agricultural work became mechanized, society didn't grind to a halt. On the contrary. While the rise of factory-style farming may have created temporary employment problems, over the long haul those who would have been farmhands were incentivized to get different jobs, including ones using new machines.

Sure, the world lost a sizable number of manually laboring hands that had been planting and harvesting on small, locally owned farms. Sustainability advocates then lamented eroding agrarian sensibilities, the diminished quality of modern agribusiness food, and the increasingly normalized maltreatment of animals. But, all this said, deploying a narrow neo-economic lens has meant focusing rather on the flipside of loss —

that is, on how well the engine of prosperity chugs along. Wages rise and the price of many desired goods drops. It's all good from this perspective.

But even if we perform the narrow economic squint, prosperity isn't guaranteed to last forever. What if the familiar two-step pattern of old-job-ending/new-job-beginning gets disrupted and the old happy endings are no more? What if the majority of workers displaced by advancing technological modes of production don't get opportunities because their hands, hearts, and minds are too slow or too costly to employ? What if meeting the demands of growth incentivizes companies to hire technological surrogates? Indeed, what if advanced machines stop needing human operators to get jobs done, and the soulless logic of capitalism fuels ever-growing inequality?

These questions may seem purely hypothetical. But Ford sees plenty of evidence that justifies soberly posing them at this particular historical moment, when "machines themselves are turning into workers, and the line between the capability of labor and capital is blurring as never before."

On the one hand, old stories, like the one about agricultural work, are getting updated. For example, the Japanese are now using devices to identify ripe strawberries by color and pluck them within a matter of seconds. On the other hand, automation isn't just taking over repetitive, assembly-style mechanical labor. In other words, it isn't just limiting options for human employment in expected sectors. Rather, Ford claims that the future of the service industry writ large might well be affected, as exemplified by Momentum Machines. This San Francisco start-up is poised to introduce cost-cutting technology designed to whip up, in rapid-fire fashion, gourmet burgers "at fast food prices." Such technology won't just supply restaurants. Momentum Machine's business plan includes stores and "perhaps even vending machines."

More shocking, Ford proclaims that even white-collar jobs are becoming precarious. Everywhere he sees signs of their impending obsolescence. *Forbes* and other venues are publishing computer-generated news. Law firms are using eDiscovery software to analyze documents. Professors are handing off essays to computers to grade. The London Symphony orchestra played a well-received musical composition created by a computer.

The technological developments driving these examples lead Ford to conclude "there are good reasons to believe that America's economic Goldilocks period has […] come to an end."

A "Goldilocks economy" exists "when growth isn't too hot, causing inflation, nor too cold, creating a recession." Like Goldilocks's porridge, it is "just right." Ford appropriates the related idea of a "Goldilocks period" from Jared Diamond's account of the 19th-century European colonization of Australian agriculture:

> Like American economists in the 1950s, the Australian settlers assumed that what they were seeing was normal, and that the conditions they observed would continue indefinitely. They invested heavily in developing farms and ranches on this seemingly fertile land.
>
> Within a decade or two, however, reality struck. The farmers found that the overall climate was actually far more arid than they were initially led to believe. They had simply had the good fortune […] to arrive during a climatic "Goldilocks period" — a sweet spot when everything happened to be just right for agriculture. Today in Australia, you can find the remnants of those ill-fated early investments: abandoned farm houses in the middle of what is essentially a desert.

Heated debate surrounds the future of the United States, and Ford marshals plenty of evidence to justify pessimism. To give two examples, he emphasizes that, "as of 2013, a typical production or nonsupervisory worker earned about 13 percent less than in 1973 […] even as productivity rose by 107 percent and the costs

of […] housing, education, and health care have soared." And the other example: "The first decade of the twenty-first century resulted in the creation of no new jobs" and "income inequality has since soared to levels not seen since 1929."

Reviewing all of Ford's statistics as well as chasing down interpretations that challenge his analyses would lead us down a well-traveled economist's rabbit hole. To get a fresh perspective, we ought instead to consider whether Ford is right in identifying a *techno-optimism bubble* that impedes smart discussion of the economic consequences of innovation.

Ford only uses the word "bubbles" a few times in his book. An important use occurs when he differentiates recurring economic issues from the stark new problems that technology poses. Consider the following paragraph:

> Among practitioners of economics and finance, there is often an almost reflexive tendency to dismiss anyone who argues that this time might be different. This is very likely the correct instinct when one is discussing those aspects of the economy that are primarily driven by human behavior and market psychology. The psychological underpinnings of the recent housing bubble and bust were almost certainly little different from those that have characterized financial crises throughout history […] It would be a mistake, however, to apply that same reasoning to the impact of advancing technology.

Here he runs the risk of conflating two different things: 1) whether something new is actually happening with technology itself; and 2) whether longstanding psychological tendencies make it hard to detect technological novelty. Indeed, if many readers find Ford's skepticism ridiculous, maybe their ridicule goes beyond standard disbelief supported by wonky economic calculations. Maybe naysayers like Ford have a hard time being taken seriously because of the "social-psychological phenomena" that give rise to contagiously optimistic bubble-thinking.

Vincent Hendricks, professor of Formal Philosophy and director of the Center for Information and Bubble Studies at the University of Copenhagen, notes: "The term 'bubble' is no longer confined to just financial movements […] [I]t can refer to irrational, collective, aggregated behaviour, beliefs, opinions or preferences based on social proof in all parts of society." According to Hendricks, "boom-thinking, group-thinking, herding [and] informational cascades" are some of the main mechanisms that lead to bubbles forming and not bursting until devastating damage occurs. Perhaps these or related mechanisms partially account for the tendency that most bothers Ford — the tendency to anchor the future in "lessons gleaned from economic history." They don't concede or recognize that "the question of whether smart machines will someday eclipse the capability of average people to perform much of the work demanded by the economy will be answered by the nature of the technology that arrives in the future."

When seen through the lens of bubble phenomena, it does indeed seem that many people may be in denial about the likely devaluation of human labor with advancing information technology. Professionals and intellectuals of various stripes are happy to concede that low-skill and low-education jobs are vulnerable to automation (and, of course, globalization), but assume their own jobs will be shielded from technological takeover. They cling to the outdated belief, shouted across the marketplace and educational sectors, that investing in higher education of an intellectually demanding sort will ensure success in those jobs that computers can't dominate.

It's an understandable conviction. If it's wrong, then today's lucrative positions, now showered with social prestige, will be eliminated. Widespread advice about how to stay a step ahead of machines will come to seem hopelessly idealistic. And cherished convictions about sophisticated human judgment being irreducible to computational processes will be nothing more than Pollyannaish dogma. Ford challenges comforting status

quo thinking on the subject by declaring:

> As the technological frontier advances, many jobs that we could today consider nonroutine, and therefore protected from automation, will eventually be pulled into the routine and predictable category. The hollowed-out middle of the already polarized job market is likely to expand as robots eat away at low-wage jobs, while increasingly intelligent algorithms threaten high-skill occupations.

Ford does concede some fields will better withstand the automation onslaught than others. For examples, he characterizes healthcare as an especially robust sector. While those "areas of medicine" that "don't require interaction with patients" (e.g., radiology) are on shaky ground, healthcare workers who physically and conversationally engage the sick are in much better shape.

In the end, however, Ford depicts resilient professions as rare. He foresees a tidal wave of apocalyptic destruction arising from potent overlaps between big data (exploding through cloud computing) and artificial intelligence (amplified by deep learning):

> organizations are collecting incomprehensible amounts of information about nearly every aspect of their operations, and a great many jobs and tasks are likely to be encapsulated in that data — waiting for the day when a smart machine learning algorithm comes along and begins schooling itself by delving into the record left by its human predecessors.

Like escaped prisoners too busy focusing on where they're going to notice the trail of clues they're leaving behind for their captors, workers everywhere, according to Ford, are creating massive digital footprints that reveal how they deliberatively and intuitively solve problems. Ever-smarter machines will follow their tracks in order to turn the data into reliable, computational heuristics and problem-solving techniques.

And that, in a nutshell, is the root of Ford's diagnosis of the problem. We're human. In the absence of specialized training that pushes against our natural inclinations, we're biased to see the world through human eyes. From such a perspective, today's accomplishments will be met or exceeded by future generations of humans. But the world looks different from a machine's perspective. The machine's focal point will be massive databases and other digital repositories overflowing with enough raw material to create templates of skill that don't require humans for their operation.

Ford's bubble hypothesis suggests we're too optimistic about the future for three reasons: we've got all-too-human sensibilities; we're biased to believe experts, including economists, who channel such sensibilities; and there's little incentive to depart from status quo thinking because, if we do, then we're left with only two rational options — to pursue healthcare work; or, against the odds, to become precisely those disruptive entrepreneurs creating technologies that speed up our collective impotence. These aren't great choices for a lot of us. As for shifting to a collective action solution, like demanding a guaranteed income for all, this seems inconceivable in the American political context.

One way to assess the bubble hypothesis is to see how it stands up when applied to specific cases. Let's take the future of law. In "Four Futures of Legal Automation," Frank Pasquale and Glyn Cashwell essentially pump the brakes on enthusiasm for automating legal work; they note that while some parts of the litigation process have been delegated to machines, more complex areas of the law — many of which, it seems, cannot be fully understood without direct experience or rich sociological understanding of their operations — are less amenable to computerization. Highlighting the differences, they caution against prognosticators who extrapolate from cherry-picked cases (supporting the "humans-out-of-the-loop" model) to the field in general:

> Classic values of administrative procedure, such as due process, are not easily coded into software language. Many automated implementations of social welfare programs, ranging from state emergency assistance to Affordable Care Act exchanges, have resulted in erroneous denials of benefits, lengthy delays, and troubling outcomes. Financial engineers may quantify risks in ever more precise ways for compliance purposes, but their models have also led to financial instability and even financial crisis.

Pasquale and Cashwell also argue that "difficult cases" that go beyond "settled law," such as deciding how to handle emergencies like cybersecurity breaches, might strain the ability even of highly advanced artificial intelligence. In these instances, competing value-based judgments would need to be made, and computers might lack the tacit knowledge and normative sensibilities to make those judgments.

And then there's the diagnostic problem. When automation skeptics point out problems with existing systems, they're essentially providing uncompensated labor that helps software companies remedy mistakes. This criticism-response dynamic can become a perverse feedback loop that funnels into automation advocates and scholars claiming "progress." It might have disincentived Pasquale and Cashwell from disclosing specific problems they've encountered with automated legal software.

Finally, there's the political problem. Pasquale and Cashwell argue that in the years ahead it will take more than "the degree to which tasks are simple or complex" to determine which legal jobs get assigned to humans and which to machines. Sociological and political factors will also play a role, with "extralegal developments" proving "crucial" for "determining the future balance of computational and human intelligence in the law." For example, if political will grows for combatting inequality, then society may determine that "human judgment" is needed to execute the critical legal tasks associated with reviewing complex regulations and drafting measures for better regulating markets and sectors that challenge social justice ideals.

These issues — selective choice of examples, misjudging the computer's perspective, market complications holding back automation critics, and automation speculators' tendency to divorce technological options from political decisions — apply to prognostication in lots of fields, not just in the law. As the debate about automation continues, it will become increasingly important to weigh them against the biases that concern Ford.

And that's why it's not enough to fight predictions with counter-predictions. Whether we live in a bubble that prevents us from looking into the abyss of future unemployment depends in part on the psychological, social, political, economic, and even technological conditions that shape the marketplace of opinions on the matter. Futurism needs to be tempered by rigorous inquiry into its own underlying dynamics.

¤

Evan Selinger is a professor of philosophy at Rochester Institute of Technology, where he is also affiliated with the Center for Media, Arts, Games, Interaction, and Creativity (MAGIC).

13

Jailbreaking Thought: The Alan Turing Centenary

by Eli MacKinnon (September 30, 2012)

Eli MacKinnon uses the occasion of the centenary edition of Andrew Hodges's 1983 biography, Alan Turing: The Enigma*, to remind us that, for the father of computer science, the essential value of the field lay in its potential to elucidate the secrets of human thought.*

WHILE OUT FOR A RUN, Alan Turing conceived of the computer. He wasn't mid-stride. By his own recollection, he was taking a break, lying in a meadow a few miles south of King's College.

He saw that any mathematical operation that could be performed by the teeming neurons of a human brain could also be worked out by a single machine with only three capabilities: distinguishing between two symbols, traversing a series of those symbols and swapping one for the other. With this insight, he had invented a method by which to offload humanity's entire wealth of formalized logic onto a discrete cluster of inorganic matter. As he saw it, he'd found a way to share thought itself with the bloodless parts of the universe.

Andrew Hodges writes in his still-definitive 1983 biography, *Alan Turing: The Enigma*: "His machines — soon to be called Turing machines — offered a bridge, a connection between abstract symbols, and the physical world." Hodges' book was reissued in honor of the June 23, 2012, centenary of Alan Mathison Turing's birth, a marker that touched off a variety of celebrations and reassessments of the marathon-running, code-breaking, man-loving war hero's life. Among the most apt of these was an interactive Google doodle that offered a very close replica of a Turing machine in its original conception.
Hodges addresses Google's debt to Turing in his new preface to the book:

> Online search engines, which work with such astonishing speed and power, are algorithms, and so equivalent to Turing machines. They are also descendants of the particular algorithms, using sophisticated logic, statistics and parallel processing, that Turing expertly pioneered for Enigma-breaking. These were search engines for the keys to the Reich.

The Enigma referred to here is the machine that the Germans used to encrypt military transmissions during WWII — a kind of modified typewriter that would bounce a letter through an ever-shifting series of circuits until it flashed onto a lampboard as a distantly abstracted cipher. Germany's confidence in the impenetrability of its messages never faltered. In terms of immediate impact on the lives of "pink-colored collections of sense-data" (Turing's half-joking phrase), it was his work on cracking and re-cracking the Enigma that represents

Turing's greatest achievement. Now, no one doubts that his invention of a machine to undo the Enigma — to rapidly dispatch vast swaths of the German device's quintillions of possible configurations — saved Allied lives and drastically altered the course of WWII. But during Turing's lifetime, the history-altering work he performed for His Majesty's Government was a secret to the public, a fact that made possible the many tragedies of his final years.

Pointing to a 2011 speech in which Barack Obama mentioned Turing in the same breath as Darwin and Newton, Hodges acknowledges in the new preface that "public recognition of Alan Turing [has] attained a level much higher than in 1983, when this book first appeared." He's right, but given the man's impact on nearly every aspect of modern life, his profile is still mind-bogglingly low. As Hodges makes clear, Turing was the antithesis of a self-promoter, stubbornly believing that the merit of his thought alone should always be enough. This disregard for publicity, as well as the fact that his most sensationally patriotic work wasn't revealed to the public until its urgency had passed, can explain some of the discrepancy between his significance and his renown.

But there's another fact that may account for much of our cultural amnesia — or indeed cultural suppression from the outset — regarding the man who humiliated Nazis while running marathons at an Olympic level and then found time to do what amounted to early R&D for our iPhones. Turing was conspicuously homosexual, and he couldn't be bothered to hide this fact at a time when his sexuality was illegal. The punishment for his openness – a compulsory hormone regimen that caused him to grow breasts — arrived when he was 39, administered by the government he helped save. Not long after, Turing killed himself.

It's fortunate that Hodges had the passion and foresight to write his record when so many who knew Alan Turing were still alive to help decipher his life. Hodges' enthusiasm for his subject animates all 540 pages. An English mathematician himself, he never winces at a technicality, but he's at his most gripping and expressive when describing the man's tics and attitudes, the "voice that went up and down in pitch rather than in stress," the ink-drenched collars.

Perhaps this is why the opening section on Turing's childhood, whose intimacy and narrative cohesion is never again matched, remains the book's most haunting. It gives early glimpses of his genius for abstraction: at 10 years old, he complained that during an introductory lesson to elementary algebra, his teacher "gave a quite false impression of what is meant by x." And it makes good on the always-satisfying historical shaming of naysayers. After catching him doing algebra during "religious instruction," a teacher wrote,

> I can forgive his writing, though it is the worst I have ever seen, and I try to view tolerantly his inexactitude and slipshod, dirty, work, inconsistent though such inexactitude is in a utilitarian, but I cannot forgive the stupidity of his attitude towards sane discussion on the New Testament.

But the chapter's core is the devastating story of Turing's first love, a boy named Christopher Morcom who bonded with Alan over science at boarding school. A few gentle warnings from Morcom that Turing was showing him too much attention ensured that their relations never strayed far from the intellectual. The boys tracked comets together by post when they were separated during school breaks, and at the time of Morcom's sudden and unexpected death from bovine tuberculosis, their friendship had only been deepening. In Hodges' telling, the loss seems to crystallize Turing's future, imbuing him with a sense of solemnity, and a drive to compensate for Morcom's stifled brilliance.

The recollections that Turing wrote soon after his friend's death are among the most moving passages in the book; his plainly expressed longing for Morcom reverberates when we see him drifting from one unrequited love and rebuffed advance to the next as an adult. In a postscript to the condolence letter he sent to Morcom's mother, Turing couldn't help asking for a photograph: "I should be extremely grateful if you could find me

sometime a little snap-shot of Chris, to remind me of his example and of his efforts to make me careful and neat. I shall miss his face so, and the way he used to smile at me sideways. Fortunately I have kept all his letters."

By the time WWII arrives and Alan agrees for the first time to "surrender a part of his mind, with a promise to keep the government's secrets," as Hodges puts it, the biography has shifted to a tone that is part political thriller and part math text.

Readers who take the time to internalize the technical details of Enigma-breaking will be rewarded with visceral insights into the kind of lateral thinking that can occasionally place a single, peculiar Englishman at the tiller of world history. Though some of these sections may drag for readers with no special interest in cryptography or math, it would be a mistake to fade away before learning that Turing asked for a teddy bear for Christmas months before he prefigured the computer age with his epiphany in the meadow; that he was moved to sponsor the boarding-school education of a young Jewish war refugee from Vienna, whom he proceeded to make an unsuccessful pass at; or that somewhere in the backwoods of Buckinghamshire are two bars of silver bullion that he buried as an insurance policy against a German invasion, encrypting instructions on their retrieval that mysteriously failed to turn up the treasure after the war.

The Enigma's second half confronts us first with the slow-rolling bureaucratic indignities that Turing endured as he tried to realize, in collaboration with the British government, the first working version of the "universal machine" he dreamed up on his run, and then with the senseless cruelty of the punishment he suffered at the hands of the state for "gross indecency." He earned this charge after a brief and perhaps naively optimistic (or perhaps just unlucky) affair with Arnold Murray, a working-class youth he met on the streets of Manchester. A friend of Murray's who knew about the affair and understood that sex criminals didn't enjoy the usual protections of the law saw an opportunity to burglarize an upper-middle-class man's house. The thief wasn't excessively greedy, but when Turing realized what had happened, he made the mistake of acting as an innocent man. He called the police and then quickly incriminated himself in describing the likely culprit as a friend of his boyfriend's.

Turing went on to choose the absurd pseudoscience of a two-year estrogen treatment, which smothered his libido, but not his sexuality, over jail time. At the time of his death on June 7, 1954, the treatment had been over for a year, and in the weeks leading up to that day, he had been working on new ideas, socializing within his small group of friends and in most other ways behaving in his usual good-humored manner. This information, as well as the fact that he was an increasingly visible homosexual in possession of important state secrets, has given rise to speculation that Turing's death may not have been a suicide. But in Hodges' account, the official judgment of a June 10 inquest — that Turing had poisoned himself with cyanide, apparently delivered in a partly eaten apple found by his bedside — emerges as the likely reality, an outcome motivated by outside forces but not decided by them (Hodges does allow that the suicide may have been mercifully orchestrated to allow one person, Alan's mother, to believe that it was an accident).
"I am forced to the conclusion that this was a deliberate act. In a man of his type, one never knows what his mental processes are going to do next," wrote the coroner. And for his own reasons, Turing may have agreed with him. After all, during a 1952 BBC radio panel discussion on whether machines could think, he had said, "Thinking is those mental processes we don't understand."

Exposition of Turing's uncompromising philosophy on the nature of thought and his foundational work in the field of artificial intelligence makes up the brighter spots of the book's second half. The metaphysical provocations of the Turing Test, an exhilaratingly simple probe of machine intelligence, have only become more urgent since he proposed it in 1950. The test asserts that the best way to determine if a computer is really thinking is to ask it yourself, over and over again and in as many ways as you can devise, until either you

satisfy your skepticism or the machine reveals itself as a hollow pretender to a human right.

Now that mass-marketed cell phones can survive a few rounds of the question game, it's easy to lose sight of the fact that, for the father of computer science, the essential value of the field lay not in its everyday utility, but in its potential to elucidate the secrets of human thought. Turing once wrote, "I am more interested in the possibility of producing models of the action of the brain than in the practical applications of computing." It is perhaps this purity of intellectual motivation — a refusal to retreat from the controversial to the marketable or from the truly mystifying implication to the merely useful application — that represents Turing's most vital characteristic. Hodges' contribution is to convey that fearlessness to future thinkers of all makes.

¤

Eli MacKinnon is based in China; he works in Shenzhen's tech industry and follows its young startup scene. Previously he worked as a reporter for Live Science.

14

Shaping the Future of War: The Important Questions

by Julie Carpenter (May 23, 2016)

Julie Carpenter reviews two books — Christopher Coker's nonfiction Future War *and P. W. Singer and August Cole's science fiction* Ghost Fleet: A Novel of the Next World War. *Her review focuses on how "ethics" struggles to keep up with technology, particularly during war.*

In 1911, Italian pilot Giulio Gavotti became the first person to launch an aerial bombardment. En route to a Turkish camp in his monoplane, he had the novel idea of throwing grenades over Libya. It was a momentous decision — his action would spawn the era of modern aerial warfare with its attendant strategies and counterstrategies. That there was a clear-cut tactical advantage for the side with the most Gavottis (by way of pilots, planes, and explosives) was uncontroversial. But, significantly, as with most new tactical advantages, aerial bombardment sparked ethical concerns — about technological and economic disparities in warfare, for example. These ethical niceties were anything but clear-cut or uncontentious.

To this day, we can see echoes of such concerns — about what's fair in war. Do pilots, like Gavotti, have an advantage considered inhumane, even in the scope of war, because of the distance between killer and killed? What exactly is "inhumane" for that matter in the context of war? Do today's Remotely Piloted Aircraft (RPA) personnel, or *drone* operators, experience greater emotional removal from war atrocities because of their increased distance from the destruction they cause? Or do their prolonged surveillance activities and extensive use of drone cameras bring them closer to the people on the ground emotionally, even though the pilots themselves are physically removed from danger? Do their emotions even matter? More pragmatically, do such technologies save lives by removing pilots from harm's way? And, more generally, is it the nature of the tools used, and the ensuing carnage, to jump-start new debates about ethics and technology?

There are certainly no easy or static answers to these questions, but two recent books — Christopher Coker's nonfiction work, *Future War*, and P. W. Singer and August Cole's science fiction story, *Ghost Fleet* — provide thought-provoking, albeit very different, frameworks for thinking about them. They help us to predict, or at least imagine, the likely future of war. They also help us understand the history of that future.

¤

"War doesn't change," declares Coker, a professor of international relations at the London School of Economics and Political Science. Rather, what changes "is the way we see it — perception is everything." How we

interpret its causes, sacrifices, economic or personal gains, participatory roles, or any aspect of war, morphs over time. If war was once, and still is in some places, an occasion for the exuberant celebration of manhood or nationhood, it is now more likely viewed as a necessary evil whose harms must be mitigated. But, however it is regarded, how we determine what tools to use, and what policies to develop for their use, is shaped by a wider culture than the military or even militias. Again, the current use of aerial force operated by pilots at a greater distance than in the past has not only elicited much public debate, but sparked new policy efforts at the UN level by such NGO groups as the Campaign to Stop Killer Robots. These are responses rooted in a specific cultural context. As Coker points out in his book, while the emphasis in these discussions at first appears to be rooted in new technologies and the novel ways we use current technologies, the essence of what changes is not, in fact, our technologies or even our access to them. What changes is the way we must assess the world and how we act out in it and with it. As Georgetown University professor David Cole recently observed, "Technology cannot solve the moral and ethical issues; it only casts them into sharper relief."

In *Future War*, Coker does not claim to present a volume of "lessons learned" applicable to future technologies. Instead, he brandishes an analytical tool he terms "future thinking." "Just as there is a visitable past, there is also […] a visitable future, a world whose consciousness is one we can imagine because it may not be entirely dissimilar from our own." By means of nuanced and multidisciplinary interpretations of significant points in the history of warfare, he extrapolates forward. For instance, he turns to the example of technological interactions among technologies that no longer have humans "in the loop" as one "historical change that allows us to visit the future." In machine-to-machine communication over the internet, the preponderance of web activity is currently nonhuman, with things like search engine optimization bots and data scrapers wading through more content than people ever could. Such interconnected technologies will proliferate to ever more domains of life — we will, of course, quite happily learn to live with technologies that, to use a quotidian example, allow our coffee maker to tell the web-based coffee store we need another whole bean delivery of our preferred brand by same-day drone-to-door transport. That interconnectivity has military applications we need to explore.

In terms of how we communicate with machines, that, too, will continue to evolve in ways we need to explore and extrapolate from. Coker uses the example of search engines, which we originally turned to for alerts on what was "important." We now turn to them for what is "true" — thus, as a primary source of knowledge, which again has relevance to how we think about, even conduct, war. To Google the word "war," Coker points out, means the search engine will direct us to the sites that it thinks are most relevant or illustrative of our information goal. This relevancy is determined through algorithms that use, for instance, general popularity as a guide. No longer is the machine just mediating human-human interaction: what's added is a very complicated layer of machine interpretations to the answers it provides to the knowledge we seek.

To take this idea a step further, we can look to the example of the robotic exoskeletons that soldiers or first responders may soon be wearing. They are projected not only to enhance human strength and endurance, but also to monitor and dynamically promote the wearer's health by, for instance, triaging bleeding wounds with a styptic foam applied to bodily injuries that are sensed by the suit itself — without purposeful or direct human input. Part of the communication between human and machine will be at this "unconscious" level, occurring below the threshold of human intention or knowledge. How will that "unconscious" communication evolve in the context of war? Or of peace for that matter?

Once we begin to imagine the spectrum of near-possible technologies and their uses, policy and philosophical questions further proliferate. Sticking with the robotic exoskeleton example: Will the technology be available only to military personnel, or to everyone? If it's available to the public at large, and surely it will be, then will the suit's cost be so prohibitive that it further separates the economic haves from the have-nots? Will "the haves" become so enhanced that different rules apply to them? Or, is another set of questions more relevant:

Will collecting continuous biodata via the exoskeleton constitute an invasion of the wearer's privacy? How will that medical data be used and protected? Who is responsible for the safety of the wearer in different situations, the individual or the exoskeleton manufacturer?

It goes without saying that we can expect many more such revolutionary technologies in the near future, which means it is imperative that we begin the important discussions around their policy, ethical, and cultural impacts now in order to more thoughtfully integrate these innovations into our everyday lives, whether at home or in battle spaces. The coffee example is not banal, and we shouldn't imagine it is. The looming privacy issues surrounding the Internet of Things is not banal either. Both will shape and be shaped by militarized spaces. The fact is, though, that at present, the policy-making process is enormously backlogged: it is playing catch up to a host of paradigm-shifting technologies — those already mentioned above and others like, for instance, Remotely Piloted Vehicles (RPV), domestically used quadcopters, and DIY 3D-printed weapons to name just a few. The timeframes we have historically devoted to generating ethical and philosophical frameworks regarding safety, privacy, potential off-label uses, are no longer available to us. We have to find ways of generating policy more quickly and nimbly.

The goal of future thinking is, then, to fill this gap, between what we are inventing now and what we will live with tomorrow. It is to think ahead, rather than play catch up. Coker's intention is, thus, pragmatic: to "future think" in order to imagine and perhaps implement interventions sensitive to cultural context. In a nutshell, it is to enact practical change ahead of actual invention and proliferation.

¤

An engaging aspect of Coker's manuscript is his use of science fiction metaphors to frame and, just as importantly, to make accessible complex scenarios around emerging technologies. Of course this aspect only works if you have at least a passing familiarity with such popular culture icons as *Star Trek*, William Gibson, Ray Bradbury, and Jules Verne. Singer and Cole, the authors of *Ghost Fleet*, also use science fiction references liberally throughout their story, including, again, healthy references to *Star Trek*, which is held up as an example of positive future-forward thinking. This device makes sense in both books; science fiction can function as an imaginative tool for understanding how people might interact with emerging technologies. In *Ghost Fleet*, using science fiction in this way has the added effect of creating a kind of "nerd comradery" between the authors and their readers, a coded way of winking: "We know you know this reference because we're betting that you, like us, are into this stuff." Yes, Singer and Cole, we are.

In the fictional *Ghost Fleet*, Singer and Cole deploy an action-packed set of story lines to paint a picture of a new geopolitical structure. (In full disclosure, P. W. Singer and I have a collegial acquaintanceship, supported by our connections via social media and [I believe] a mutual interest and respect for each other's real-world work on humans interacting with technology in defense spaces.)

Writing about the future of anything is a bold, daunting, and thought-provoking task, and Singer and Cole's real-life think tank background is clearly to their advantage here as authors — Singer as a strategist and senior fellow at the New America Foundation and Cole as a nonresident senior fellow at the Brent Scowcroft Center on International Security at the Atlantic Council. The pictures they paint are rooted in present-day objects or artifacts. For example, the widespread use of *viz-glasses* in the book is traceable to Google Glass, and, as readers, we can clearly recognize how the geopolitical landscape described in *Ghost Fleet* could happen, too.

Singer and Cole play out their next-world power battles in a variety of locations — in cyberspace and outer space; in the forests and beaches of Hawaii and in the cityscapes of Beijing; and in the eponymous floating graveyard of an abandoned Naval fleet. These ships — the same "ghost" ships from of the book's title — are

abandoned, floating, reserve Navy ships left by previous generations as placeholders of war. In a word, they are obsolete. Without dipping into spoiler territory, the ships are resurrected to act again as instruments of war. The book's missions and battles illustrate why the touchstones of history and human innovations should not be so easily dismissed as obsolete, nor should they be disentangled from the future: different forms of fighting have their strengths — from advanced technologies that are semi-autonomous, to older technologies that call for a human-in-the-loop. To be sure, the newer ships might have "better" machinery with more sophisticated weaponry than this ghost fleet, but that very technological "advancement" relies to an enormous extent on the infrastructure in place, and so can be just as vulnerable, albeit in different ways.

Consequently, the technology is a complicated presence in itself, hard to pin down as something external, or Other, because it is so fully integrated into people's bodies; it is a part of their very selves at a biological level, a part of the human way of developing, living, and being. Throughout the book, people on board are, in philosophical terms and through literal references, considered to be "part of the ship," meaning their lives are inextricably entwined with shipboard activities, even beyond their assigned duties — although some sailors may point out that this has always been the case (think what is expected of a captain when his or her boat sinks). Being "part of the ship" also hints at what the act of framing personnel — people — to mirror the fate of a ghost fleet might mean. If our very human bodies with all their frailties are boosted with *viz-glasses* and *stims* (synthetic stimulants), we are in essence upgrading our biological systems, but to what end? And at what cost? Will just being human ever be enough again, or will people who do not integrate technologies into their physical selves be left behind as archaic or ineffective?

Ghost Fleet imagines Chinese forces taking advantage of their American counterparts and launching a historically familiar surprise attack on the US naval base at Pearl Harbor. Triggered by a Chinese gas discovery near the Mariana Trench, the Chinese government, known here as the Directorate, has enough power and economic leverage to act without fear of repercussions (in the form of sanctions) from the United States. The battles that follow feature the likely future military integration of humans with weaponry, by way of autonomous drones, university-supported cyber militias, and individual biological boosts via *stims*, that increase soldiers' reaction times and alertness — the latter ingested almost constantly by those who rely on their effectiveness.

Cultural acceptance of, and widespread reliance on, tools like *viz-glasses* divides the military culture even at a team level, amplifying and complicating divisions in the ranks based on age and experience. To illustrate a classic trope: The older fighters are reluctant to embrace the so-called advances exemplified by these new gadgets, while the younger ones use these technologies as part of their standard military toolkit. Although the majority of fighting actually takes place on traditional land and sea, Singer and Cole use their professional expertise to imagine, or "future think," warfare in space and cyberspace.

The story of this new "future war" is revealed through the perspectives of multiple people around the world. Nonstop action paints a cinematic picture that does not let up, inviting easy comparisons to Tom Clancy's works and occasionally to such storytelling cousins as *Marathon Man* and *A Clockwork Orange*. One caveat: Sometimes these contexts overshadow the multiple story lines and individual perspectives. In other words, this is a book driven by action more than character or emotions. Stories are painted with broad brushstrokes — characters sometimes fall into distracting potholes. A "Black Widow" assassin and the Chinese Directorate's Admiral Wang are film noir-ish characters, both of them largely defined by their obsessions: her personal mission to seduce and kill Chinese soldiers in Hawaii in Bond-like fashion, and Wang's love for relying on his encyclopedic knowledge of War Philosopher-General Sun Tzu in almost every exchange with other characters.

I expect *Ghost Fleet* will eventually be the basis for a movie. It is easy to imagine the multiple plots' relentless action lending itself to on-screen representation. Coker's *Future War* is also energetic in tone, and intellectually

intriguing in how it melds geopolitics, history, defense, cultural change, and technology. The same recurring motifs animate *Ghost Fleet*: increasing physical and mental reliance on technologies, the ambiguities of war, and the dynamic effects of rapidly changing cultures.

In both books, we should, as citizens of the world, recognize ourselves as characters in the stories of politics and war. *Future War* shows us paths we have taken and suggests where we are heading. *Ghost Fleet* imagines a world of people where humans remain not only relevant, but their very flaws and humanness make for dependable heroes and antiheroes when the call of duty — and storytelling — arises. All this said, we are left with more questions than answers. For instance: Will integrating systems like immersive, web-connected augmented reality glasses with our bodies make us less human or superhuman? Will humans remain invested in wars for duty, profit, and passion? And, the singularity question plaguing so many of us: Will the mechanical minds of AI continue apace until they are told to stop, or will they spontaneously develop their own set of internal guidelines and motivations?

¤

Julie Carpenter's research focuses on human-robot interactions. Her 2016 book on the topic is entitled Culture and Human-Robot Interaction in Militarized Spaces: A War Story. *She is currently a Research Fellow in the Ethics + Emerging Sciences Group at California Polytechnic State University.*