

The Transparency Paradox at Work

By Anna A. Tavis

This Perspective shines new light on the evolution of the transparency and privacy debate in the age of artificial intelligence (AI), also known as “machine learning.” Whether we call it “people analytics,” “HR metrics,” or any other data-related name, we are still talking about human behavior—recorded, aggregated, broken down, and served back to us as data. With the ubiquity of mobile devices and the always-on Internet, the moment of truth comes with the question of who is seeing our data and how it is being used. Take HR or management, what data should be used to recruit, appraise, promote, or dismiss employees? Take employees, how much of our private lives do we want our employers to know about and judge us by?

Ethan Bernstein opens the debate noting that with machine learning the very notion of “transparency” in management has taken on a new paradoxical dimension. “People say they value their privacy, but they don’t act that way.” He draws a distinction between “data fusion” that feels more like surveillance and top-down control and “personal data fusion,” which is about employees using their own data to improve their performance. The real ethical question for 21st century management becomes whether managers, peers, or HR teams can actually decide how much data they are willing to “not access for the sake of productivity.”

The Hogan Assessments team offers their psychological angle on the debate. Online behavior is still fundamentally human behavior, they argue, and “the difference between the 20th and 21st centuries lies in the ubiquity and volume of behavior,” broad access to data, and “the sophistication of tools to interpret the results.” The authors argue in favor of human nature that is “inherently social.” They are optimistic that humans will continue to “swim best in a world



of connection, relationships, status hierarchies, and groups.” The introduction of increasingly more powerful tools will help us tame technology in the service of even greater humanity.

Doug Cunningham, a developer who runs his own HR technology business, knows first-hand the power of data and the possibilities it offers. He cautions on the efficacy of the decisions that could be made deploying AI. We should not get carried away with what is technically possible, but rather focus on the boundaries that need to be placed on the data used. Cunningham calls on managers and HR professionals to join in national and global community conversations to collectively anticipate the future evolution of AI and be prepared to address the consequences.

Robin D. Richard adds to the exchange a compelling illustration of how the solution to the transparency paradox could be found in the human-centered design. Take the case of his company, CareerArc, which

provides personalized and on-demand outplacement services correcting for the fatal flaws of the brick-and-mortar outplacement firms, lack of privacy, and availability of choice. With the ubiquitous technology available to us, the service is scalable and democratic, addressing the need of the new economy for the ongoing career development options. The solution to the transparency paradox will be possible only when “products are designed to respect, not correct, the human instincts for privacy and safety.”

The final word in this discussion comes from **Christine Congdon** of Steelcase, an innovator in smart workspace design. Steelcase is optimistic about technology and sees it as an enabler that will take human performance to the next level of productivity and excellence. Taking its cue from the latest car design, she argues that if cars could actually help people be better drivers, then tech-enabled offices will help people be more productive and more connected workers.

This Perspectives comes full circle to conclude that technology is on its way to help humans be better humans, depending on how we use it and why

Anna A. Tavis, Ph.D., is Perspectives editor and an associate professor at New York University. You can connect with her on LinkedIn at [annatavis](#).

POINT

The Evolution of “Transparency” in Management

Get Me Everything You’ve Got on ... Me

By Ethan Bernstein

It’s startling—but no longer science fiction—how much information can be and is being gathered about employees at work. Fifty years ago, a typical manager might have periodically tracked a few numbers—revenue, expenses, customer satisfaction. Today, the workplace is bristling with monitoring software, sensors, and cameras. We label our workplaces “smart” because they are always observing us: a cocktail of smartphones, computers, fixtures with embedded sensors, and cameras collectively contain enough locational, audio, video, text, and activity data to produce an unfathomable set of digital breadcrumbs. And the more information management gets, the more it realizes it needs.

More Is More

More information not only calls for even *more* information, it also demands more information-processing power. No person or HR team could ever parse so much data, but people no longer have to. That’s why we have artificial intelligence (AI), or machine learning. AI can filter floods of information—from our email, apps, calendars, social media, Web browsers, news services, enterprise workflow apps, systems of record, monitoring devices, wearable sensors, video camera feeds—and *make sense of it*. All

in real time. While we humans can only handle so much data, AI systems get smarter with more information.

Our work lives are now full of examples. Google can use your past location and calendar data to predict your next client visit and help you avoid traffic jams on a trip you haven’t even told it you will be making. Siri can answer your questions with much greater ease and accuracy because she knows everything you have ever done with—or near—your iPhone (and its apps, microphone, accelerometer, GPS sensor, Bluetooth, WiFi, and so on), and when she hears you curse or sees you frown, she can ask you if she can help. Cortana can tell you to “ask me anything” because, chances are, the answer lies somewhere in the gigabytes of information flowing through—or stored in—your work computer. Even your office bathroom’s hand soap dispenser (enabled with RFID sensors which read your ID badge) can remind you, by name, to wash your hands before you return to work. This is not science fiction. Even in ordinary workplaces, substantially increased use of observation—“the act of careful watching and listening, or paying close attention to someone or something, in order to get information” (as the Merriam-Webster Online Dictionary defines it)—has become widespread over the last 15 years.

As I explain in my forthcoming article in the *Annals of Management*, these Big Data developments, while viewed as cutting-edge, are in fact the latest phase of a long evolution of observation in management. It’s pretty basic: you can’t understand, control, or change anything without observing what’s going on. During the Renaissance, it was an innovation to keep modern accounting records. Roughly a century ago, it was an innovation to observe scientifically exactly how workers did their work. A few decades later, it was an innovation to observe scientifically how workers worked together. Now it’s AI-enabled “transparency.” With each innovation, companies have learned more about how they could improve productivity—if they collected the necessary information.

Which Side Are You On?

Studying this history, I noticed that

there has long been something one-sided about it. It’s all conceived and implemented from the observer’s point of view: *We* need to know what *you* are doing. The more we know, the more we can help you do it better. But what about the observed? Are they willing to be so intensely examined? Does it actually help them do better? Or do they now clamor for “privacy”?

Management—including HR—has largely been able to brush these questions aside for two reasons. The law allows them to collect as much information as they can, as long as they let employees know they are doing it. And employees are letting them do it (just click “Agree”). This is actually somewhat puzzling. In what social scientists call the “privacy paradox,” people say they value their privacy, but they don’t act that way. Think of it. Most of us, even if the offer were made, would not sell our private information to a big, rich company for a paltry \$60. Yet Microsoft recently bought LinkedIn and what it bought—at \$60 per person—was all that information that 433 million of us put onto our LinkedIn pages. The same thing happened when Facebook bought WhatsApp for all its address books, and Facebook got away with paying only \$42 for each one. Not that any of us got a cent. If we value our privacy, why did we put all that information online for free, knowing that companies sell it to each other?

Even so, companies aren’t necessarily getting the free ride it appears they’re getting. It has long been known—and we know from our own experience—that human behavior may change when we know we’re being watched. As I explained in my 2014 *Harvard Business Review* article, “The Transparency Trap,” when we feel over-observed at work, our performance suffers. This can take two forms. One response is to just do exactly what the watchers want to see. Observers may get compliance, but they won’t get much innovation. We’re just not likely to try something different if we’re being watched to make sure we’re doing our jobs right.

Another response, no more to a company’s advantage, is to find ways to hide. Put employees in open offices

and they'll work from home—and feel more productive. Track more data, and they'll find a way to stay under the radar. Monitor their work smart-phones, and they'll get a second one, as many people do. Make all written work accessible, and people will stop writing things down. Track email, and they'll use Slack instead. These are all real examples. So rather than learning more, management may end up learning less, or even learning things that aren't true.

Let “Us” Have the Data and You Can Have the Results

Must all this incredible data-gathering and data-crunching ability go to waste, then? I don't think so. It just needs to be rechanneled. When entrepreneur David Brunner was conducting his doctoral research at Harvard, he observed that organizations can deploy AI either to help them improve their employees or to help their employees improve themselves. The first approach is standard *data fusion* and to employees, it feels like surveillance and control (even if you might call it “people analytics” or “transparency”). The second approach is called *personal data fusion* and to employees, it feels like coaching, mentoring, and self-improvement. It's for *them*. The company provides AI to gather and process information which *you*, the employee, own and which *you* can use to improve *your own* performance. You become a more valuable employee—which of course can be rewarded by the organization. You are also likely to feel more loyal. Thus, the company gets the results it wanted, not by taking, but by helping its employees give.

In that sense, AI presents an amazing opportunity. Whereas being observed by a person can feel like an invasion of privacy, it doesn't if the “observer” is a machine—the data stays personal so long as another *human* doesn't access it. And if the machine adds value to how we do our work (like with LinkedIn and WhatsApp), we give willingly. But at work, this only works if the data stays personal—i.e., out of the curious hands of another human being. Yet putting together a business case for privacy is far from easy in a world blindly enamored with transparency. So

whether you are manager, peer, or HR team, the real question about transparency in today's workplace is how much data are you willing to not access for the sake of productivity?

Ethan Bernstein, Ph.D., is assistant professor at Harvard Business School and author of *The Transparency Trap*. You can find him on Twitter @ethanbernstein.

Data Fusion Is Unlikely to End Life as We Know It

By Dave Winsborough,
Darko Lovric, and
Tomas Chamorro-Premuzic

A few years ago, when Cambridge University researchers demonstrated that your Facebook likes could be used to predict your sexual orientation, voting preferences, personality, and even IQ, the world erupted in a frenzy of outrage and concern, and governments everywhere responded with new laws limiting the use and sale of your private data.

No they didn't. Nothing much happened at all. There were news stories, mild wonderment and a smattering of online chat, but the net effect on online behavior was precisely zero. Facebook continued its stratospheric growth, and while a few thousand people may have tightened their privacy settings the net impact of knowing that Facebook could know everything about you was a collective, “meh.”

Most users of social media acknowledge that there is a tradeoff between what you give (access to lots of personal data) and what you get (free search, a free platform to follow celebrities, and the largely unfulfilled promise of relevant ads).

Nor do they seem to mind that these data are used to evaluate them. For example, employers and recruiters use social media to evaluate job candidates, and at least one business, LinkedIn, was designed and built to capture user profiles to monetize recruitment and search at a global scale. Our own research and that of others shows that younger people expect that employers peruse their online

lives and are comfortable with being evaluated in that way. Employers regard it as a method for “discovering the applicant's true self.” After all, online behavior is still behavior and people take as many pains online to curate an image of themselves as they do offline. When social media users decide what images, achievements, musical preferences, and conversations to display online, the same self-presentation-dynamics are at play as in any traditional social setting.

Burnishing one's image online has even been taken as a right, insofar as the European Union allows citizens to hide links to images or posts that do not fit the reputation they want to portray online. Consequently, people's online reputation

The fragmentation of tools to merge these data, and the changing fashions between applications and formats means that there isn't a master algorithm that binds our fragmentary digital traces into a coherent whole, a digital twin, who contains our thoughts, emotions, and actions.

is no more “real” than their analogue reputation; the same individual differences are manifested in virtual and physical environments, albeit in seemingly different ways. It is therefore naïve to expect online profiles to be more genuine than resumes, although they may offer a much wider set of behavioral samples.

The difference between the 20th and 21st centuries lies in the ubiquity and volume of behavior that is captured, by the breadth of those with access to those data, and by the sophistication of tools to interpret the results. Right now, the fragmentation of tools to merge these data, and the changing fashions between applications and formats means that there isn't a master algorithm that binds our fragmentary digital traces into a coherent

whole, a digital twin, who contains our thoughts, emotions and actions.

You can bet good money that in the future our digital selves will have personalities that are accessible to anyone who cares to look. It will be, possibly, an age of true digital transparency.

Yet when Socrates exhorted us to “know thyself,” it is doubtful he considered a world where companies (read your employer, or Facebook, or Google) know us better than we can know ourselves. That opens two scenarios for the future.

In the first we lack agency over our digital selves. Our personal data is owned instead by the hardware and software we use and sold to the highest bidder. Anonymity and privacy are in demand but very expensive—but there is no practical opt-out for ordinary people. In this world our every engagement with the digital world creates food for marketing and social engineering of Machiavellian/Orwellian proportions.

In the second scenario, our personal data is ours. New companies have emerged to allow us to “bank” our data and “lend” it to corporations. These companies have the opportunity to allow our digital selves to become our agents and avatars, revealing our preferences for specific purposes. We’ll have as many digital selves as required, ensuring fragmentation works for us instead of against us. Importantly, these companies can use our profiles for both self-insight and growth, helping us lead more authentic lives but preventing others from using this information to manipulate or game us.

Humans are inherently social, and we swim best in a world of connection, relationships, status hierarchies and groups. Technologies that emerge from human activity are still fundamentally human. Through history, social technologies have been decried as destructive (cellphones were predicted to destroy face-to-face interaction; rock and roll would usher in the end of days; and video games would stop children from exercising). Personal data fusion is very likely to happen. The world is very unlikely to end.

Dave Winsborough is vice president of innovation at Hogan Assessment Systems and founder of the New Zealand firm Winsborough Limited. He can be reached at dave@winsborough.co.nz.

Darko Lovric is a principal at Incandescent and a former fellow of the World Economic Forum. He can be reached at darko.lovric@incandescent.com.

Tomas Chamorro-Premuzic is the CEO of Hogan Assessment Systems, a professor of business psychology at University College London, and a faculty member at Columbia University. He can be reached at tomas@hoganassessments.com.

Applying AI

By Doug Cunningham

The business world is buzzing with visions and promises of how artificial intelligence will radically alter companies. The potential is real. The number of workplace data sources is growing at a breakneck pace, and so is our ability to store, process, and draw conclusions from this data.

The result: an unprecedented opportunity for AI to uncover patterns and associations that yield perspective and insights on our workforce. Areas such as employee engagement, workplace health, talent acquisition, and productivity are all positioned for dramatic change.

Tread Carefully

We have already witnessed rapid developments in the consumer technology market. Apple’s Siri, Google Now, and Microsoft’s Cortana are becoming smarter every year. Most of us welcome the advancements and gladly yield our privacy in exchange for the benefits.

But applying AI in the workplace exposes companies to a different set of risks, and leaders must take great care to protect their employees and businesses. Companies have a responsibility to protect the privacy of their employees and ensure they are not discriminated against.

Privacy. When sufficiently powerful AI crunches vast quantities of data on web searches, emails, chat messages, calendar events, mobile activity and more, it has the potential to learn personal information that employees never intended to share with their employer. One could argue that employees should

take care in using company systems for personal use. But anyone who has worked in a corporate setting knows that consistently separating work and personal is near impossible.

Is it acceptable if AI learns that an employee has financial problems? Is expecting a child? Has a sick family member? With sufficient data, these aren’t difficult to figure out. A watchful manager might draw similar conclusions. AI can watch everyone all of the time.

Discrimination. The real power of AI is to do something with what we have learned, and this is where businesses need to be even more careful. Julia Angwin recently presented the results of a study analyzing the output of software in wide use across the country to predict future criminal behavior. They uncovered a racial bias that had previously gone undetected.

Businesses leveraging AI need to approach this with utmost diligence, lest a recruiting chat bot wrongly reject a specific group of applicants or a performance system unwittingly develop a gender bias.

What to Do

So how do we ensure that we don’t learn things we shouldn’t? Or make decisions we shouldn’t? Should we run from AI? Of course not. But we should be very thoughtful and considered in how we build and deploy such systems.

To start, artificially intelligent systems in the workplace must be built not only considering the possible but also the allowed. The vendors behind them must develop safeguards to monitor their results and detect breaches of privacy and unintentional discrimination.

Second, managers and executives deploying such technology need to be cognizant of the risks and how they use the technology. They need to set clear guidelines around acceptable use and communicate these to employees. They need to work closely with vendors to ensure that any technology deployed aligns with the expectations of the business.

Last, we need continued discourse on the challenges and opportunities. The White House Office of Science and Technology Policy this year hosted pub-

lic workshops to prepare for the future of artificial intelligence. Other efforts are under way as well. Management and human resources executives should participate in these groups to expand their view of the possible, but also understand the risks and concerns.

Doug Cunningham is CEO of Uppercase HR, an agile performance management software company. He can be reached at doug@uppercaser.com.

The Paradox of Innovation and the Currency of Trust

By Robin D. Richards

Bernstein's writings on privacy and personal data reminded me of the unintended paradoxes that can sometimes result from product innovation and design. Much like in his article, *Big Brother Vs. Personal Data Fusion*, where he offers the example of modern CRMs that come with the promise of delivering more useful and transparent data to sales organizations, but may in reality encouraged secrecy in the salespeople they were designed to assist—we found a similar paradox in our study on workplace flexibility.

When we surveyed over a thousand working professionals on their work-life balance, telecommuting arrangements, and wellness benefits, we found that one in three working Americans felt obligated to answer work-related calls and emails past regular working hours. With the company-issued smartphone tethering modern-day workers to the virtual desk, it appears that time saved by many employees from using the technology initially designed to make work and life more efficient, often ends up overtaken with more work. However, it is important to note that it is the person who, shaped by their culture and standards of practice, decides to rededicate those hours to work, and not the innovation itself, that causes the paradox.

As human capital software and service providers, we at CareerArc think about these types of decisions every

day. Developing technology for people who manage teams, recruit talent, and transition employees, we build each new product platform or feature with a people-centered approach to design. An area we have applied this focus is one where technology actually *supports* the human instinct for privacy and safety, and that is in our approach to modern outplacement.

Many of today's outplacement services, which are benefit given to laid off employees to support their job search and transition, still follow the brick-and-mortar model born out of the 1960s: They typically include the provision of physical satellite offices where coaches conduct resume and interview preparation in-person. But for companies downsizing, this model can be costly and increasingly underutilized, effectively reserving outplacement for executives and upper-level employees and excluding those in entry-level to middle-management who could arguably most benefit from assistance. We immediately saw the need for a tech solution that could impact more people.

Designing with our end-user—the exiting employee—in mind, we aimed to correct the two design flaws in the brick-and-mortar approach that were similar to the flaws Bernstein found in many data and enterprise tools today: the neglect of the human instincts for safety and privacy, and the omission of choice (forced opt-in, and so forth).

Knowing job loss and career transition can be one of the most challenging moments in a person's life, we learned through customer interviews and feedback that privacy, safety, and autonomy were very important for those undergoing this change. By delivering career assistance content and virtual coaching through an online platform, we strived to create a safe space for these new job seekers to brush up on their resume writing and interviewing skills, take self-assessment tests, explore new career interests, and even video chat with a coach all in the privacy and convenience of their home. Most importantly, users are given full autonomy over their time and priorities; they decide what to do first, when to do it, and whether to opt-in at all. Through this approach, we are seeing employees land jobs three

times faster than the national average time it takes to find employment.

Whether it's moving from an offline note-taking system to the sales CRM in the cloud, or transitioning from face-to-face outplacement models to on-demand video sessions with a career coach, the core issue here is trust—the real currency in today's increasingly con-

The core issue here is *trust*—the real currency in today's increasingly connected and surveilled work life.

nected and surveilled work life. Personal data fusion, a system that can strike that perfect balance between data privacy and transparency, is possible only when products are designed to respect, not correct, human instincts for privacy and safety.

Robin D. Richards serves as chairman and CEO of CareerArc, a leading HR technology company powering employer branding solutions that help business leaders recruit and transition the modern workforce. He can be reached at rrichards@careerarc.com.

Technology Drives the Well-Being of People

By Christine Congdon

Offices would be better places to work if they could learn from cars. New car models are embedded with technologies that make driving easier, safer, and more fun. Sensors tell drivers if there is a truck in their blind spot, or if they are about to back into another car when parking. The car doesn't just provide transportation anymore—it actually helps people be better drivers.

People used to think that technology would make offices obsolete—but the opposite is happening. In the near future, technology will be embedded in offices so it actually helps people work better and makes the workplace even



more relevant. Like cars, a network of sensors and other technologies in the workplace will help make work a much better—and more humane—experience. Technology will serve individual workers, teams, and organizations. It will help people cope with the sense of being overwhelmed that they often feel as work intensifies and the pace of change accelerates. This data will also help organizations design the kinds of spaces that workers love to work in versus have to work in.

Companies that want to create great workplaces can benefit from this embedded technology to help individual workers and teams, and they can also draw from the data that is generated. Design, facilities, and real estate professionals can make better decisions about where to focus their efforts if they have a data stream to tell them which rooms are always busy and which rooms are ignored. With this information, orga-

nizations can better understand what's working and what's not, so they can

The challenge with technology in the workplace is making it meaningful to the employee.

make the best workspace possible.

The challenge with technology in the workplace is making it meaningful to the employee. It should help them cognitively off-load some of the tasks they have to think about today, and leverage new technologies that will be embedded in the physical environment to make their work lives easier and more productive. When smart and connected spaces are enabled in the office in an intentional way, people can more easily navigate the complexity of work

today and reduce their stress.

When objects, like chairs or rooms, can sense the environment and communicate, they become tools for understanding complexity, identifying opportunity and responding to needs swiftly. The work environment itself will become a tool for creating more productive, engaged employees who are in control of their surroundings and able to choose what they need to accomplish their tasks.

Just as technology in today's cars is improving the driving experience, tomorrow's office will harness the power of emerging technologies and allow people to more easily navigate the complexity of work as well as help organizations create better work experiences for individuals and teams. ■■

Christine Congdon is global director of research communications at Steelcase. She can be found on Twitter @cscongdon.