

# Predicting the unpredictable: How will technology change the future of work?

*AI is having its moment in the workplace. But what does the future hold for worker and AI collaboration? It depends on the decisions we make and our ability to see beyond them.*

ARTICLE • 20-MIN READ • 26 JUNE 2024 • Center for the Edge

The future is an unexplored frontier, and our intuitions about what we'll find there are often wrong. Christopher Columbus famously thought that he was establishing a new trade route to the Indies when he stumbled across the Americas. Similarly, predictions about the future of work typically bear little resemblance to the actual future once we find ourselves standing in it. Consider economist John Maynard Keynes' 1930s estimate that a 15-hour work week would be possible within a few generations.<sup>1</sup> Technological advancements, he predicted, would improve productivity to the point that humans could enjoy the same living standards with less time devoted to work. Historian and author Rutger Bregman reused Keynes' prediction in 2017,<sup>2</sup> as it remained unrealized. Bregman estimated that a 15-hour work week should be possible by 2030. But despite some early gains in the modern era,<sup>3</sup> the typical work week since the 1940s seems to be stuck at a little under eight hours a day, five days a week.<sup>4</sup>

Why do our predictions so often turn out to be wrong? For business leaders, inaccurate predictions can become a problem as they lead to poor decision-making and missed opportunities. Focusing only on trends limits our understanding of the decisions in front of us. When predicting how technology (like generative artificial intelligence) will transform the future of work, and thus better prepare and equip the workplace for its impact, it is important for leaders to be able to see *all* possibilities and possible futures on the road ahead—as well as *all* of the decision points that will determine which possibilities and futures become reality.

## **Deloitte Insights Magazine**

Read more from issue 33 of Deloitte Insights Magazine on advancing the AI conversation



## **The problem with predictions**

The real problem with our predictions isn't that our technique is poor or imprecise, or that we overestimate the potential

of new ideas and technologies in the short term while underestimating their impact in the long term.<sup>5</sup> We could easily rectify these issues by improving our forecasting processes, by tapping into more diverse data sources or integrating superior technology (quantum computing, perhaps) to increase the quality of our predictions. Our challenge is different. Our challenge is to rethink our prediction models in ways that allow us to see the broader spectrum of options and opportunities.

A forecast relies on model, a way of framing the present and its dynamics. The model identifies the nouns (the actors involved) and, more importantly, the verbs (the interactions between actors, and between the actors and the environment). The model we use determines which paths to the future we can see, which possibilities and opportunities are visible to us, and which are invisible or hidden. Different models enable us to see different possibilities, and possible futures.

The models we use struggle to account for humans and all their desires and inventiveness. The future is shaped by a myriad of human decisions and it's these human decisions that determine which future we find ourselves in. Discounting (or implicitly not considering) this human factor means that our predictions assume that society is heading in one direction when society often decides to make a sharp turn and go in an entirely different one.

This is perhaps most apparent in how our adoption of new technologies plays out. How we choose to use technology is as, if not more, important, than the technology itself.<sup>6</sup> New technology creates new possibilities, but it's up to us to determine which possibilities crystallize into actualities.

The rush to automate the home in the mid-1900s is a great example. All manner of labor-saving devices were developed to support the homemaker. *Design for Dreaming*, a promotional film from 1956,<sup>7</sup> suggested that we were heading toward a push-button future where automation would free homemakers from the drudgery of maintaining a home. This is not the future that we live in, though. Instead, automation enabled society to raise our standards and expectations for a well-maintained home, and the work required to maintain a home became more onerous as a result.<sup>8</sup>

We can see this dynamic in operation today with concerns that new workplace monitoring technologies—such as facial recognition or artificial intelligence-powered performance management tools—are driving us toward a workplace dystopia.<sup>9</sup> This is not an inevitable future, though—only one potential path we could head down. If we do find ourselves in a future dystopia, it won't be because technology forced us there. Instead, the future we find ourselves in will be the result of many small (and possibly well-meaning) decisions on how technology was integrated into organizations' performance management processes.<sup>10</sup> It is more productive to consider technology as a form of human action, rather than as some extrinsic force we have little control over.

The challenge when trying to understand the future is that it's difficult, if not impossible, to see past decisions we don't understand. Moreover, we are often unaware of which decisions will be the consequential ones at the time we're making (so many of) them. We might not even realize that there is a decision to be made—a choice—because the model we're using to frame the present doesn't allow us to see the choices before us. If we want to peer into the future, then we need to reconsider the models we're using, to look for complimentary or alternative models that enable us to see different futures, different possibilities, and different opportunities.

## Rethinking our framework for 'work'

How we frame the present—and the challenges we have today—determines how we think about the future. Was the shift in work at the start of the 2020 COVID-19 pandemic a shift from working in an office to working remotely—a change in where workers are located? Or was it a shift from working physically to working digitally—a change in where the work exists, a change in work medium? Physical work is dominated by pen and paper, physical tools and materials, and in-person interactions, making it necessary for workers to gather, usually in an office or a factory, to be productive. Digital work, on the other hand, relies on digital documents, media, and interactions, freeing workers from the need to be in the same physical location as their materials, tools, and co-workers.

---

# It is our imagination (or lack of it) that determines our ability to access the new opportunities and possibilities of working digitally.

---

Consider a self-driving bus that needs a human driver to attend an accident it is involved in. The work needs to be done where the bus is, so it's clearly *somewhere* work and so toward the bottom of the grid. It's also important that the driver is certified to drive the bus (should they need to manually drive the bus past some obstacle), though it's not important which driver in particular attends. This would put it midway across the horizontal axis—not anyone can do the work, but there are potentially a few workers who could. We might compare this to a chief financial officer signing off on their organization's books. Clearly there is only one person who is allowed to do this (the CFO) making it *someone* work, but the work could be done digitally and so located anywhere—placing it at the upper-left extreme of the grid. Or we might decide to hold a collaborative team meeting in person (rather than digitally) to foster intra-team relationships. This makes the meeting *somewhere* work, though the team could have a great deal of freedom in the *where*. Someone is clearly defined though—the members of the team.

If work is physical, then it is more productive to gather workers in physical workplaces. We might find, for example, that the creativity and innovation that we value so much decline when workers work remotely.<sup>11</sup> If this is the case, then working remotely, where workers rarely (if ever) head into an office, is likely a temporary diversion. Competition in the workplace will soon drive workers back to the physical workplace where they are more productive.

These effects, however, might be an artifact of legacy work practices that rely on in-person interaction. It's true that early experience with remote collaboration in the 1990s showed lower levels of creativity than when collaborating in person. But more recent experience has remote collaboration yielding better results than working in person, the result of improvements in remote work technology and work practices, making physical colocation less important.<sup>12</sup>

Physical work is becoming more amenable to being done digitally as technology and work practices improve. Instrumenting a forklift by remote control,<sup>13</sup> for example, frees the forklift driver to work from where it is most convenient, rather than where the forklift is located. This enables us to rethink how manual and knowledge workers collaborate across the (digital) workplace, developing new and more productive work practices. Modern remote work can encompass much more than the knowledge work we traditionally associate with it. It is our [imagination](#) (or lack of it) that determines our ability to access the new opportunities and possibilities of working digitally.<sup>14</sup>

Working physically and working digitally provide us with different *affordances*. An affordance refers to the perceived or actual properties of an object or environment that suggest how it can be used. A door handle, for example, affords pushing or pulling based on its design, and this design provides users with a cue about the intended action. An affordance concerns the relationship between an individual and an option, highlighting the possible actions that the user perceives the object to have. Those affordances present us with different opportunities, and so force us to make different cost-benefit trade-offs. Physical work, for example, requires us to warehouse workers in offices and factories. Current work practices have evolved to take advantage of this colocation. Digital work, on the other hand, makes a distinction between where the worker works (a physical place) and where the work lives (in the network of digital media that binds a team together). Place becomes the free variable when working digitally as workers are no longer required to be in the same location to work together. This can have a detrimental effect on current work practices, such as making some conceptual tasks—developing new ideas and designing research, for example—more challenging.<sup>15</sup> On the other hand, working digitally makes it easier to pull together a diverse team across geographies and tap into particular expertise—two strategies that can improve team performance.

This doesn't imply that we can't achieve the same outcomes while working digitally as when working physically. Before we can realize the benefits of digital work, new practices that make the most of the advantages of working digitally need to be developed. Developing these new practices will require a great deal of experimentation and learning by doing. The first

step is to recognize that it is how we think about work that is holding us back.

Consider how we have historically framed work: skilled workers in a physical workplace. We focus on how technology can automate tasks and consequently drive skill churn. The future is quantified via estimates of which skills (and so workers and jobs) emerging technologies will make redundant. Recent estimates of the impact of large language models (LLMs) by the International Labour Organization, for example, estimate that 24% of clerical tasks should be considered highly exposed to automation effects with an additional 58% having medium-level exposure,<sup>16</sup> while another report suggests that 80% of US workers could see at least 10% of their tasks automated.<sup>17</sup> The International Monetary Fund estimates that AI would affect 40% of jobs and worsen inequality,<sup>18</sup> while Goldman Sachs predicts that up to 25% of the work currently done by humans could soon be automated by generative AI.<sup>19</sup>

These predictions assume that automation displaces skills, forcing workers into a skill upgrading cycle where they supplement lower skills that have been automated with higher, yet-to-be-automated skills (typically 21st century skills such as creativity, critical thinking, collaboration, information literacy, etc.)<sup>20</sup> via retraining. But how do we know that this skill displacement model accurately captures the consequences of introducing new technology into the workplace? And that reskilling and retraining are the best responses for workers?

Hiding behind this focus on tasks and skills is a model of how work is done: product-process-task-skill, the skill-based division of labor. Skilled workers complete specialized tasks within formalized processes that result in the creation of products. Our focus is on the individual workers and the attributes which enable them to perform the tasks presented to them.

The product-process-task-skill model tells us that introducing automation into a workplace drives skill upgrading. Our first (simple) attempts at automation target lower, simpler (and so easier to mechanize) tasks. Workers are driven to higher tasks (and so skills), unable to compete with the more productive (and tireless) machines. Automation targets increasingly advanced tasks as it is developed and becomes more capable, driving workers to increasingly higher-level skills. Eventually workers are forced out of their jobs, when the technology becomes more competent than the humans, with industries “de-skilled” as skilled labor has been replaced by (unskilled) machines. Telephone exchange operators are a classic example of this dynamic. Automation was first used in simpler local exchanges, driving workers to learn to manage more complex regional exchanges. Improvements in the technology then enabled automation of regional exchanges, driving workers to upgrade their skills again, to manage international exchanges. Eventually all exchanges were automated, displacing workers entirely. Much of the anxiety about the current AI moment is that AI seems to be capable of performing more and more high-level tasks, whereas early waves of automation were mainly concerned with lower manual skills and tasks.

Framing our analysis differently—considering a different model—enables us to unlock new possible futures, possibly more desirable futures.

One possibility is to consider the work *system*, rather than the *individual* when trying to understand work. For example, is the task of flying a commercial airliner conducted by a pilot alone? Or by a system composed of two pilots along with cockpit instrumentation and controls (and possibly even a flight engineer)?<sup>21</sup> We might even extend this system to include air traffic control towers (along with the controllers working in them), flow control networks, and other factors and organizations that influence the work of flying a plane.

Introducing new technology—automation—into this system can lead to a range of effects. We can think of each effect as a pathway to different possible futures. Skill upgrading with consequent de-skilling is one possible pathway, where human pilots are replaced by increasingly more capable robotic doppelgängers, eventually resulting in an autonomous plane. Another possible pathway is skill downgrading, where automation displaces higher-level skills, driving the pilots to lower skills. Our doppelgänger might monopolize the serious flying, relegating pilots to driving planes from the end of the runway to the boarding gate (something we might now do remotely, much like the forklifts mentioned earlier).

We should also consider lateral pathways, such as how the introduction of new technology into the cockpit results in expertise being redistributed within the work system. We might introduce speed reference cards and throttle indicators to support the process of coordinating airspeeds with takeoff and landing wing configurations for aircraft of different

weights. Remembering throttle positions is no longer something that the pilot must do, as the cockpit instruments transform the pilot's task from remembering critical speeds to making judgments that integrate nuanced environmental factors.

Considering the system, rather than the individual, enables us to consider worker attributes that don't fit into our current conception of skill. Creativity is a prime example, as we now understand that creativity is the result of a generative process:<sup>22</sup> Interactions in, and influenced by, the workplace build on domain knowledge, past experience, and differing perspectives on the problem at hand to synthesize a novel and useful response. Creativity is something we do in a place, rather than something we have in ourselves. These non-skill attributes also include those that are not associated with particular tasks, such as a worker's situational awareness. Toyota, for example, realized that automating a factory and eliminating workers entirely improved production in the short term but reduced longer-term productivity as there were no workers embedded in production to discover (and point out) changes that might improve the process. The solution was to reintroduce workers, giving them roles in production that valued them not for their ability to perform tasks, but for their ability to support longer-term productivity improvement.<sup>23</sup>

## New models for human and tech collaboration in the workplace

If we want to consider alternative futures, to see new possibilities and opportunities, then we need to develop new models for how work changes when new technology is introduced.

Consider the recent waves of AI-powered solutions that have entered the workplace. It's commonly assumed that AI is a task automation technology, with the main difference between this and previous generations of automation being that AI enables us to automate higher-level knowledge work tasks, rather than just lower-level manual ones.

While we can approach AI this way, it may be more productive to frame AI as automating behaviors, rather than tasks. AI enables us to codify decisions via algorithms: Which chess piece should be moved? What products are best to populate this investment portfolio? These decisions are made in response to a changing environment—our chess opponent's move, a change in an individual's circumstances. In fact, often it's the environmental change that prompts the action. This responsiveness to external stimuli is why it's more natural to think of AI as mechanizing behaviors rather than skills. This is true even for LLMs. We might prompt an LLM to *remember* something under the assumption that training has caused it to record some fact into its trillions of weights. This is not the case, though. Rather than being *recalled*, the memory is *created* when a prompt interacts with the LLM's language prediction processes. The memory is in the prompt (the words) as much as it is in the weights, much as how a smell (our prompt) can evoke a memory (a recreation) of a long-forgotten moment.

A behavior is a reaction to the world changing around us, something we do in response to external stimulus. The same stimulus in different contexts triggers different behaviors. LLMs are a case in point: Subtle changes in the prompt we provide deliver quite different results, in the same way subtle changes in how we train the LLM can change its response in surprising ways.

If we're to understand the possibilities of AI-powered automation, then we should focus on how introducing automated behaviors will redistribute expertise within the workplace, rather than the tasks and skills we might associate it with.

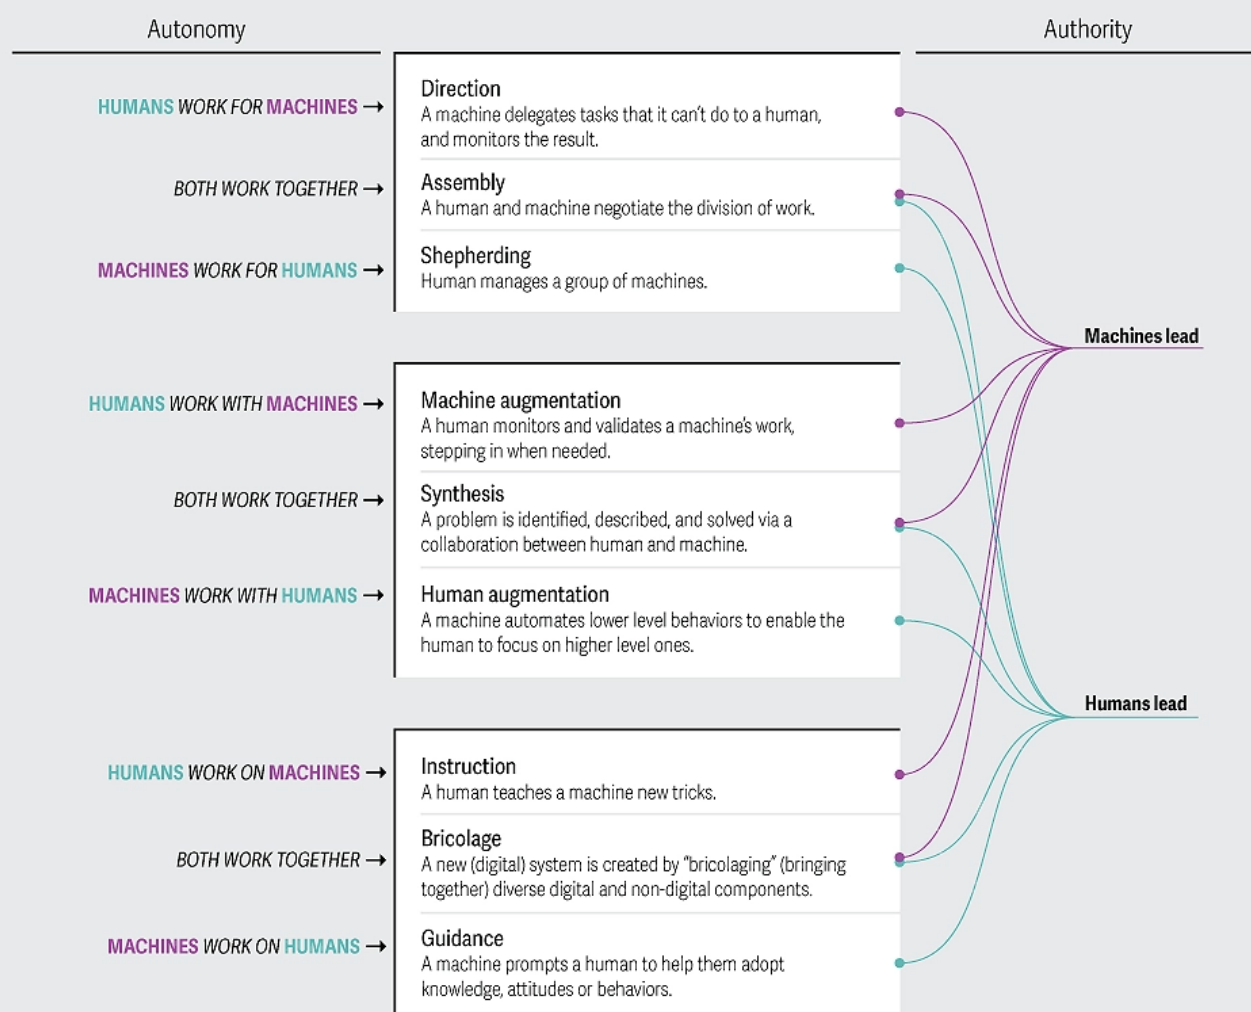
If we're automating behaviors, then there are two factors we should consider. First is the ability of the behavior to effect change, its freedom to act—the behavior's *agency*. Agency is a question of degree, rather than being binary, as it depends on both the capabilities of the automated behavior and any limitations we place on it. A self-driving taxi might be quite capable of finding its own way down an uncrowded street but falters when it encounters a problem it cannot navigate its way around and calls for a human.<sup>24</sup> Authority, our second factor, represents the superior-subordinate aspect of the human-machine relationship: who has final decision rights. Should an algorithm (AI) decide to initiate recovery for supposed overpayments to welfare recipients? Or is human oversight required?<sup>25</sup>

Both dimensions of this model (figure 1) represent a division in responsibilities—a redistribution of expertise—between human and machine. The left side represents agency: who does the work. Does the machine work on the human's behalf,

or the human on the machine's? Or is the split more nuanced, somewhere between these two extremes? Authority is captured on the right. Does the human have final decision rights, or is the machine leading the human?

Figure 1

## Agency and authority in human-artificial intelligence collaboration



Sources: Deloitte analysis; Adapted from figure 2, Evans-Greenwood, Peter, Robert Hillard, and Alan Marshall, "The new division of labor: On our evolving relationship with technology," *Deloitte Insights*, April 9, 2019.

**Deloitte** Insights | [deloitte.com/insights](https://deloitte.com/insights)

Within this matrix we can see a range of possible future pathways. A truck driver, for example, might teach an autonomous truck how to park in a particular loading bay by manually guiding it in the first time. A robot chef learns to cook a meal by observing, and then copying, the actions of a human.<sup>26</sup> A tumor-identification behavior can augment an oncologist's ability to diagnose skin cancer by helping them locate potential tumors for investigation. We integrate AI into the work system by developing a set of integrated human and machine behaviors.

Individuals and organizations are already exploring these potential pathways, developing solutions, creating new work practices, weighing benefits against problems and limitations, and making the many human decisions that determine which future we'll work toward. Reoccurring solution patterns and relationships are already emerging (figure 2).



Figure 2

## Emerging relationships in human-artificial intelligence collaboration



Sources: Deloitte analysis; Cantrell, Sue, Thomas H. Davenport, Steve Hatfield, and Brad Kreit. "Strengthening the bonds of human and machine collaboration." *Deloitte Insights*, November 22, 2022.

Consider the following examples of how we are seeing these human/tech collaborations play out in the workplace:

- **The supervisor:** An algorithm allocates tasks—for example, a ride-sharing company that uses AI to dispatch rides to drivers who have a few seconds to accept or reject a ride request without knowing the destination or fare. Performance and pay are determined by AI. An AI also decides when morale-boosting motivational messages are needed.
- **The prioritizer:** An AI algorithm addresses a list of tasks—sales leads to be pursued, medical problems to solve, fundraising opportunities to follow up on—and ranks them in terms of their importance or potential value. The human worker then pursues them in order, sometimes with suggestions from the machine about how to do so.
- **The personal coach:** AI discovers the human worker's strengths and opportunities for improvement on a specific task (such as a telephone or video sales call), resulting in continuous engagement with AI to improve the human's performance.
- **The muse:** Multiple creative suggestions are prompted by a human, output by a machine, and iterated in an ongoing collaboration. Examples include design suggestions based on architect prompts and AI-driven generative design.
- **The collaborative decision-maker:** Complex decisions, such as medical diagnoses, are made in a dialogue between AI

and humans, where AI can improve decisions by enumerating available options, helping people weigh them objectively, and suggesting the highest probability of successful action.

- **First pass at a task:** A machine performs the first pass at a task—a life insurance application, a medical coding categorization, an analysis of an MRI scan—and makes a preliminary decision judgement. The human reviews the analysis and determines if it is correct. The order of this sequence could also be reversed.
- **The triage nurse:** AI assesses the problem (medical symptoms, for example) and decides whether a human consultation is necessary; if not, it dispenses advice to address the relatively minor problem.
- **The doppelgänger:** Machines learn from a human or group of humans to mimic their behaviors and decisions, so that the human(s) can be replicated digitally.
- **The subordinate:** AI systems perform menial, structured tasks (like extracting key data from documents or faxes) under human supervision and review.

It is unlikely that one all-encompassing model will be enough. The human-AI model we’ve just considered can help us see how the introduction of automated behaviors enables us to refactor authority and agency to create new opportunities. It has, however, little to tell us about the opportunities (and challenges) we’ll confront working digitally versus working in person.

In that instance we might want to think in terms of how some work requires the worker to be in a particular physical place—*where* the work is done is important—while other work doesn’t. Or how some work can require specific skills or certifications, or a particular reporting line—*who* does the work is important.<sup>27</sup> We can create another model, shown in figure 3. *Where* the work is done is on the vertical axis, while *who* does the work is on the horizontal.

Figure 3

### A model for digital working

	SOMEONE	ANYONE
ANYWHERE	<b>Worker-centric work</b> The work requires a particular skill or authority, and so a particular worker.	<b>Pure digital work</b> Both worker(s) and work are virtual beings.
SOMEWHERE	<b>Pure analog work</b> Both worker(s) and work are physical, and are best located in one place.	<b>Place-centric work</b> The work must be done in a particular place.

Source: Deloitte analysis.

**Deloitte** Insights | [deloitte.com/insights](https://deloitte.com/insights)

## The art and science of seeing around corners

Predicting the future is hard. It was widely assumed, for example, that the last wave of AI would make radiologists redundant, but radiology is booming as a profession.<sup>28</sup> Nor was it the case that the mass adoption of computers by business forced all workers to become coders.<sup>29</sup> And everyone got travel agents wrong, including former US President Barack Obama, who once suggested in an interview that travel agents would soon be extinct.<sup>30</sup>



Workplaces have also been transformed by technology without it ever finding its way onto the disruption radar. Sign writers and sign writing are an example. The profession reinvented itself over the 1990s as desktop publishing and computer-controlled vinyl cutters transformed sign writing from a paint-and-brush trade to design-print-and-stick. There was no industrywide disruption though. The transformation lowered costs which expanded the addressable market—sign writers went down market and also found new opportunities (such as wrapping cars).<sup>31</sup>

Our prediction track record is, at the very least, poor, and consequently, predictions on how the current crop of technologies will affect workers and workplace should be taken with a grain of salt.

Predicting the future is also something of a fool's game as we don't need to predict the future to productively engage with it. Our future is shaped by a myriad of human decisions, and it's these decisions (not the predictions) we should be engaging with—decisions that determine the future we're heading into. New technology provides us with new possibilities and opportunities, and it's how we choose to find and use these opportunities that determines the future we'll find ourselves in. Workplace surveillance is a case in point. New surveillance technology is creating opportunities to harvest detailed information on worker activities. This information presents us with new possibilities, creating an inflection point. Established workplace trends could continue, subjecting workers to increasingly granular command and control regimes that reduce job quality by increasing work intensity and reducing worker autonomy. Or we might choose to head in a new direction and establish a new trend.

To consider the longer term, we need to find and then see past the decisions that new technology creates. This means creating more and better models of how new (and old) technology interacts with complex systems like the workplace. The models we use determine which possible paths forward we can see, the possibilities both visible and invisible to us, with different models enabling us to see different possible futures. More and better models enable us to see more and better possible futures. Poor models limit our understanding of possible futures. If we're to understand the potential impact of introducing new technology into the workplace, then our first consideration should be the models we're using to frame the problem. If we're not aware of these futures, then we're not aware of the decisions we're implicitly making.

We often assume that we need to skate to where the puck will be<sup>32</sup> if we're to be successful in the longer term, when our main concern should be heading in the same direction.<sup>33</sup>

Research psychologists have been trying to untangle how humans successfully find their way through a complex and ever-changing world. Not long ago they made a surprising discovery: We don't make decisions purely in our heads but by interacting with the environment around us.<sup>34</sup> Or, put another way, we *act to decide* rather than *deciding to act*—rather than analyze, predict, and plan, we observe, evaluate (often conflicting) possibilities, and respond, fostering optionality. We don't catch a flying baseball, the canonical example for prospective control,<sup>35</sup> by predicting where it will land. Instead, we continually adjust our movement relative to the baseball so that we're heading at the same direction, while avoiding obstacles. Should we run behind or in front of second base? Is diving to catch the ball a possibility? The player strives to keep multiple possibilities alive until one of them becomes so attractive that they are compelled to commit.

The challenge is that it is difficult to see past current trends—past the hype—to understand what's driving them. As we asked earlier, did the transition to remote work during the pandemic represent a change in the place that work was done? Or did it represent a change in work medium, digital versus physical? Practically it was both: While digital technology enables us to frame the shift as a change in medium, entrenched work practices provided a strong connection back to work as a place.

To see past current trends, we need to foster optionality, accepting that there are always multiple interpretations for current events and that a surprising new technology can be applied in multiple ways. Our first instincts of how a technology should be applied are often not accurate. The electrification of factories is a good example. The initial attraction of electrification was that drawing power from the new electricity distribution networks would be a cheaper and cleaner alternative to on-premises coal-fired steam engines. Factories that made the transition realized savings of roughly 20% to 60% on fuel costs.<sup>36</sup> The bulk of the benefits of electrification came 30 years later though, when manufacturing engineers realized that distributing electric power within a factory was much more flexible than distributing mechanical

power.<sup>37</sup> Factories were reorganized, using the same production machinery and floor space but arranging the machines to optimize workflow rather than (steam) power distribution, realizing a 20% increase in total productivity.<sup>38</sup>

Rather than placing one big bet on what we think will be the winner—Betamax or VHS?—and going all in, we can make many small bets with the intention of learning about and fostering potential future choices. These small bets can be approached as real options: an economically valuable right to take up, or abandon, some option in the future. The goal is to engage a diverse range of experiences and points of view to develop a rich vein of institutional knowledge—knowledge that includes different ways of framing the opportunities and challenges considered. It’s not enough for an organization to contain experts. Leaders and decision-makers across the organization need to appreciate that multiple futures are possible and consider how these futures interact with the organization’s values and goals.

## Deliberate decision-making is what will define the future of work

The poor productivity return on technology over the past 50 years is a frequent discussion topic. Inventions such as the steam engine, electricity distribution, the motor car, and even water and waste plumbing, transformed the urban environment and the world. Recent innovations haven’t had a similar effect, nor provided an equivalent productivity boost, with recent productivity growth half of what it was during the height of the Industrial Revolution.<sup>39</sup>

One issue is that without a deliberate approach, benefits can easily be absorbed into business as usual. This is true of LLMs and generative AI—it’s not the fault of the technology but rather a lack of deliberate approach. We assume that technology is taking us somewhere, when it is really us who are deciding where we are going.

Organizations and societies need to decide how they want to reap productivity dividends. A four-day work week has always been an option. It’s just not an option we’ve decided to pursue as a society (though this may be changing). If we fail to make a deliberate attempt to develop new models and explore possible futures, then our assumptions (based on current trends) are likely to be self-fulfilling in the short to medium term but wrong in the long term.

### Continue the conversation

Meet the industry leader



**Kellie Nuttall**  
AI Institute and Strategy and Business Design  
practice leader | Partner

!

### Let’s make this work.

[Change your Analytics and performance cookie settings](#) to access this feature.

BY

**Peter Evans-Greenwood**  
Australia

**Peter Williams**  
Australia

**Kellie Nuttall**  
Australia

---

#### Endnotes

1. “But beyond this, we shall endeavor to spread the bread thin on the butter—to make what work there is still to be done to be as widely shared as possible. Three-hour shifts or a fifteen-hour week may put off the problem for a great while. For three hours a day is quite enough to satisfy the old Adam in most of us!” See John Maynard Keynes, “Economic possibilities for our grandchildren,” *Essays in Persuasion* (London: Palgrave Macmillan, 2010), pp. 321–332.

[View in Article](#)

2. Rutger Bregman, *Utopia for Realists: How We Can Build the Ideal World* (New York: Little, Brown and Company, 2017).

[View in Article](#)

3. In the late 1800s the typical work week was six ten-hour days. By 1940 this had been reduced to five eight-hour days. Subsequent reductions have involved sick days and vacations. See Dora L. Costa, “[The wage and the length of the work day: From the 1890s to 1991](#),” *Journal of Labor Economics* 18, no. 1 (2000): pp.156–181.

[View in Article](#)

4. With some variation by country. Europe, for example, chose to take a greater proportion of the productivity dividend in terms of improved job quality rather than increased wealth.

[View in Article](#)

5. Commonly known as Amara’s law as it was coined by Roy Amara, former president of Institute for the Future. Bill Gates made a similar observation in his book *The Road Ahead*: “Most people overestimate what they can achieve in a year and underestimate what they can achieve in ten years.”

[View in Article](#)

6. An observation known as Kranzberg’s fourth law, which goes as follows: “Although technology might be a prime element in many public issues, nontechnical factors take precedence in technology-policy decisions.” See Melvin Kranzberg, “[Technology and history: ‘Kranzberg’s Laws’](#),” *Technology and Culture* 27, no. 3 (1986): pp. 544–560.

[View in Article](#)

7. MPO Productions, *Design for Dreaming*, short film, Victor Solow (director), 1956.

[View in Article](#)

8. Ruth Schwartz Cowan, *More Work for Mother: The Ironies of Household Technology from the Open Hearth to the Microwave* (United States: Basic Books, 1985).

[View in Article](#)

9. Veronica Nilsson, “[Avoiding an AI dystopia for workers](#),” *International Politics and Society* (2023).

[View in Article](#)

10. Peter Evans-Greenwood, Pip Dexter, Claudia Marks, Peter Williams, and Joel Hardy, “[The trust deficit between workers and organizations isn’t personal. It’s systemic](#),” *Deloitte Insights*, August 23, 2023.

[View in Article](#)

11. Yiling Lin, Carl Benedikt Frey, and Lingfei Wu, “[Remote collaboration fuses fewer breakthrough ideas](#),” *Nature* 623, no. 7989 (2023): pp. 987–991.

[View in Article](#)

12. See Figure 3 on page 24 here: Chinchih Chen, Carl Benedikt Frey, and Giorgio Presidente, “[Disrupting science](#),” working paper no. 2022-4, Oxford Martin School, University of Oxford, April 26, 2022.

[View in Article](#)

13. Kirsten Korosec, “[Remote-controlled forklifts have arrived in France, courtesy of Phantom Auto](#),” *TechCrunch*, March 31, 2021.

[View in Article](#)

14. David Mallon, Nicole Scoble-Williams, Michael Griffiths, Sue Cantrell, and Matteo Zanza, “[What do organizations need most in a disrupted, boundaryless age? More imagination](#),” *Deloitte Insights*, February 5, 2024.

[View in Article](#)

15. Lin, Frey, and Wu, “[Remote collaboration fuses fewer breakthrough ideas](#),” pp. 987–991.

[View in Article](#)

16. Janine Berg, David Bescond, and Pawel Gmyrek, “[Generative AI and jobs: A global analysis of potential effects on job quantity and quality](#),” ILO working paper 96, International Labour Organization, 2023.

[View in Article](#)

17. Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock, “[GPTs are GPTs: An early look at the labor market impact potential of large language models](#),” working paper, arXiv, August 21, 2023.

[View in Article](#)

18. Kristalina Georgieva, “[AI will transform the global economy. Let’s make sure it benefits humanity](#),” *International Monetary Fund Blog*, January 14, 2024.

[View in Article](#)

19. Joseph Briggs and Devesh Kodnani, “[The potentially large effects of artificial intelligence on economic growth](#),” Goldman Sachs, March 26, 2023.

[View in Article](#)

20. Jenny Soffel, “[Ten 21st-century skills every student needs](#),” *World Economic Forum*, March 10, 2016.

[View in Article](#)

21. Edwin Hutchins, “[How a cockpit remembers its speeds](#),” *Cognitive Science* 19, no. 3 (1995): pp. 265–288.

[View in Article](#)

22. Research in the past few decades has shown us that creativity emerges from human interaction and collaboration. It’s a generative process: Interactions in, and influenced by, the workplace build on domain knowledge, past experience, and differing perspectives on the problem at hand to synthesize a novel and useful response. This contrasts with historical views of creativity that saw it as an attribute of a creative individual, a cognitive approach which assumes that novel ideas originate in the head. On the contrary, creativity is something we do (a verb) rather than something we have (a noun). See Ronald E. Purser and Alfonso Montuori, “[In search of creativity: Beyond individualism and collectivism](#),” presented at the Western Academy of Management Conference, Hawaii, March 2000; Rob Withagen and John van der Kamp, “[An ecological approach to creativity in making](#),” *New Ideas in Psychology* 49, 2018: pp. 1–6.

[View in Article](#)

23. Yuki Hagiwara, Ma Jie, and Craig Trudell, “[‘Gods’ edging out robots at Toyota facility](#),” *The Japan Times*, April 7, 2024.

[View in Article](#)

24. Workers from General Motors’ Cruise LLC “intervened to assist the company’s vehicles every 2.5 to five miles.” See Tripp Mickle, Cade Metz, and Yiwen Lu, “[GM’s Cruise moved fast in the driverless race. It got ugly](#),” *The New York Times*, November 3, 2023.

[View in Article](#)

25. Royal Commission into the Robodebt Scheme, [Report of the Royal Commission into the Robodebt Scheme](#), July 11, 2023.

[View in Article](#)

26. University of Cambridge, “[Robot ‘chef’ learns to recreate recipes from watching food videos](#),” *ScienceDaily* LLC, June 5, 2023.

[View in Article](#)

27. Peter Evans-Greenwood, Alex Bennett, and Sue Solly, “[Negotiating the digital-ready organization](#),” *Deloitte Insights*, March 30, 2022.

[View in Article](#)

28. Gary Smith and Jeffrey Funk, “AI has a long way to go before doctors can trust it with your life,” *Quartz*, June 4, 2021.

[View in Article](#)

29. Peter Evans-Greenwood and Tim Patston, *To Code or Not to Code, Is That the Question?*, Deloitte Center for the Edge and Geelong Grammar School, August 2017.

[View in Article](#)

30. Adam Leposa, “Did President Obama forget travel agents exist again?,” *Travel Agent Central*, September 23, 2023.

[View in Article](#)

31. David Morley, “Car wrapping: Everything you need to know,” *CarsGuide*, October 18, 2023.

[View in Article](#)

32. “Skate to where the puck is going, not to where it is,” a quote mistakenly attributed to Wayne Gretzky. The original quote, “Go to where the puck is going, not where it has been,” was provided by Wayne’s father, Walter Gretzky, for a junior hockey team.

[View in Article](#)

33. It’s not good enough just to be directionally correct in the short term, when dealing with weekly or quarterly horizons. Trends will remain an important tool for managing in the short term. In these circumstances, being directionally correct is no substitute for a good sales forecast.

[View in Article](#)

34. Andy Clark and David Chalmers, “The extended mind,” *Analysis* 58, no. 1 (1998): pp. 7–19.

[View in Article](#)

35. Andrew D. Wilson and Sabrina Golonka, “Prospective control I: The outfielder problem,” *Notes from Two Scientific Psychologists*, October 9, 2011.

[View in Article](#)

36. Shifting from steam (coal) to electricity could save a firm 20% to 60% on their power generation costs, direct savings, but these savings were dwarfed by those obtained from reorganizing production, indirect savings due to a 20% to 30% productivity improvement, while using the same floorspace, workers, machinery, and tooling. See: Warren D. Devine Jr, “From shafts to wires: Historical perspective on electrification,” *The Journal of Economic History* 43, no. 2 (1983): pp. 347–372.

[View in Article](#)

37. Electrical engines deliver similar power and torque with a physically smaller engine than is the case with steam and can tap directly into electricity distribution networks.

[View in Article](#)

38. Devine, “From shafts to wires,” pp. 347–72.

[View in Article](#)

39. Robert J. Gordon, *The Rise and Fall of American Growth: The US Standard of Living since the Civil War* (Princeton, New Jersey: Princeton University Press, 2016).

[View in Article](#)

---

#### Acknowledgments

Cover image by: **Jim Slatton**

