Deloitte.

Ethics in the age of technological disruption

A discussion paper for the 2018 True North Conference

Table of contents

Deloitte believes that over the next 25 years, Canada can become the world's best place to live and work. Without a doubt, technology—much of which we can only imagine today—will play a pivotal role in our nation's progress toward that goal.

New technologies inevitably bring new opportunities and new challenges. Today, we marvel at what artificial intelligence, machine learning, and other advanced technologies allow us to do, yet we're also concerned about the impact these same technologies are having on our workplaces, our communities, and our society overall. How do we embrace technology's positives while mitigating its less desirable impacts? How do we ensure we use technology ethically, in support of what's "good," and minimize harmful outcomes?

Governments and regulators play a crucial role in answering these questions. They can react and respond to new technologies and the disruption they create—and they already are and will continue to do so. However, it takes time for governments to react and for policy to catch up. Given the accelerating rate of technological development, they can't lead in this regard.

This is where business can—and, many would argue, must—step up. Business is at the forefront of technology-driven disruption. Companies are the innovators that create the technologies changing our world, as well as the relentless competitors that use these same technologies to eke out every competitive advantage. Business leaders and their organizations are ideally positioned to ensure the technologies they develop and use function ethically, inclusively, and equitably in pursuit of "good" outcomes for all.

We think these conversations also offer an important opportunity for Canadians— based our commitment to openness, inclusivity and collaboration—to lead. In *Ethics in the age of technological disruption*, we explore the ways in which ethics, business, and technology intersect, and offer our thoughts on how businesses and governments can work together to address and overcome the ethical challenges in the age of disruption. As a sponsor of True North, powered by Communitech, we are delighted to take part in fostering this important dialogue.

Best regards,

Peter Barr
Managing Partner, Waterloo and Western Ontario Region

Ethics in the age of technological disruption

Artificial intelligence (AI), machine learning, and other advanced technologies are integrating into our everyday lives more quickly and deeply than ever. The benefits—for businesses, governments, and citizen-consumers around the world—are enormous.

But so too are the challenges. Technology-driven disruption is happening at an unprecedented pace, reshaping the way we live, work, and govern. It's creating new jobs, yet eliminating others. It's bringing us together, and driving us apart. It's creating wealth, yet exacerbating inequality. In short, it's complicated.

Today, the ethical implications of technologies and their implementation seem more prominent than ever before in our industrial history. Yet the pace of technology-driven change is so fast that policy-makers and regulators find themselves perpetually playing catch-up. Changes to policy, regulation, and legislation take time. As a result, government can only react to new technologies and their use, not provide leadership.

So how can we ensure AI and other advanced technologies achieve the best results for Canadian business and society? How do we build effective frameworks for these technologies and their use, both in Canada and around the world? In short, what does it mean to develop technology for "good"? And should business take a leadership role in this? Must business take the lead?

These aren't easy questions to answer. In the pages that follow, we'll explore the interplay between ethics—particularly business ethics—and emerging technologies such as AI. We'll look at issues of privacy, security, consent, disruption, inclusivity, inequality, and, more broadly, the role of business and government. It's our hope this paper will help spark conversations that lead to action to tackle these important issues.

Ethics
Pursuing the good

Ethics is fundamentally concerned with the pursuit of "good" actions based on "good" decision-making—decisions and actions that lead to the least possible amount of unnecessary harm or suffering.[1] The problem, of course, is that what "good" means is highly subjective, though there are certain "goods" that are more universally accepted. Even the notion of "unnecessary harm" is subjective, shaped by various influences, including geographical, political, cultural, and personal biases. As well, we humans are unable to objectively predict which choices are "best." Making a choice always entails the risk we'll choose poorly. And in making ethical choices, we must attempt to consider past and future, accounting for historical precedent while also predicting how our choice will play out in reality.

In addition, while ethics has long been seen as a matter of individual choice, today we can't escape the fact that individual ethics is inextricable from other forms of ethics—including legal, technological, political, and business.

Business ethics requires boards and leadership teams to ensure the success of their businesses in a way that limits harmful consequences for the organization itself as well as its employees and the society of which they're a part. If they wish to be ethical—and use technology for good, for example—leaders must consider all of the consequences and impacts their behaviour might have on every sphere of ethical concern. Ethical decision-making can require short-term sacrifice for long-term, sustainable success.[2]

However, this doesn't mean leaders must only choose impossibly perfect options that never cause unintended harm. Ethical decision-making must negotiate risk, reward, and safety, and it will necessarily be shaped by the principles of the leaders, their employees, and ideally, their business processes. Thinking about business ethics as they pertain to disruptive technologies like AI can seem daunting. But given that even considering business ethics has been shown to have long-term benefits for organizations, there is clearly value in exploring the ethical implications of deploying advanced technologies.

Key questions

- What is "good"?

- If we believe our businesses and technologies must "do good," how can business leaders understand and define what "good" means?

Ethical leadership and technology

Ethical leadership on the use of AI and other advanced technologies involves recognizing their potential impact, both positive and negative.

To this end, it's important to understand how these technologies affect both individual Canadians and specific groups in our society. We need to look at the commoditization of human experience, the role of consent and privacy, the significance of inclusivity and bias, the challenge of accessibility and inequality, and the ways in which we might effectively manage these issues in the age of technological disruption. Exploring these issues can help reveal how business, working together with government, can lead the way in showing how technologies can be used for good without compromising companies' ability to increase productivity and spur innovation.

The commoditization of human experience
Advanced technologies require vast amounts of data to deliver ever-more precise insights, analyses, and predictive models. And, we're creating data at an unprecedented rate. It's estimated that humans will have generated 40 zettabytes (or about 37 trillion gigabytes) by 2020, much of it by social media users.[3] Big data, the digital capturing of our human experience, has quickly become an incredibly valuable commodity—and its use and sale are the lifeblood of the digital age. Human experience has become the raw ore for a new generation of organizations to mine for profit.

Yet, as the Facebook–Cambridge Analytica scandal has demonstrated, business, government, and society continue to struggle with the ethics of the use and sale of personal data and the impact on privacy and security. In the pursuit of using big data for good, it seems that too little attention was paid to the potential for unnecessary harm.

? Key questions

- What is the responsibility of a business to its people, its shareholders, its customers, and its communities?

- What might ethical leadership look like in practice when it comes to AI and other advanced technologies?

- How can business leaders build ethics into their organization in a way that's inclusive and continually reflective of the organization's environment, employees, and broader society?

That's not to say the use and sale of data is inherently unethical; but how it's used and sold can be, especially in today's environment in which companies often buy and use data in ways that aren't addressed in user agreements—ways customers may not have expressly consented to. Therefore, in acting ethically, businesses and other organizations need to consider how the use and sale of data will affect their customers, and acknowledge that those same customers may have a right to know how their data is being used.

Consent between humans and machines

As AI systems and other advanced technologies assume increasingly central roles in the use and exploitation of data, new questions are raised that touch on privacy, trust, consent, and choice. With the help of technology, businesses today can capture more data about consumers and their behaviours than ever before, seeing patterns and making connections that others may not.

Consumers are meant to trust these technologies— and the businesses that deploy them—with all of their personal information. And they do, especially those who use online shopping services or social media, which in turn use personal data, search histories, and more to target users with tailored advertisements and other messages. But most consumers don't realize how much personal data and information these technologies and businesses have access to, or what they've agreed to in terms of the use of that data.

Legally and ethically, these issues often boil down to a question of consent, but AI systems and other advanced technologies add a layer of complexity to the concept. Current models often rely on what consumers agree to in privacy statements and terms of service. But are they enough to ensure informed consent? Research shows that the average person does not read the terms of what they agree to share when they sign up for a service—as evidenced by an experiment where 98 percent of respondents unwittingly agreed to give up their first-born child in the Terms of Service when signing up for a fictitious social networking site.[4] Moreover, what does informed consent mean in relation to technologies that often operate with little to no transparency?

Governments have work to do to catch up with developments in technology, and this can be a slow process. However, businesses don't have to wait for legislative reform—they can take the lead. Business and government can both take steps to better protect individual privacy and create meaningful mechanisms for transparency and consent. For example, businesses could enshrine better privacy protection principles in their corporate social responsibility policies and provide consumers with the information necessary to help customers make better-informed decisions.[5] Microsoft, for example, has upgraded its engineering systems to ensure it meets data protection standards and outlined its approach to AI.[6]

Key questions

- At what point does the technology-enabled collection of mundane details about individuals' lives become too intrusive?

- What rights should consumers have regarding how their data is used?

- How do we decide what data belongs to the individual?

- How can business leaders create trust in the technologies consumers are using?

- Are businesses obligated to ensure users provide informed consent of data usage that goes beyond current practices?

- Do individuals have the right to delete their data from third party platforms?

Ultimately, we need mechanisms that enable individuals to trust that technologies are using their personal data for the stated intent. But how do we create trustworthy technologies? As technology creates an increasingly globalized world, the burden for doing so will likely fall on those developing the technologies.[7] There may be a need for systems that can better explain its operations, motivations, and needs to the individual on a case-by-case basis in order to instill confidence and trust.[8]

Worker displacement and transitions

Much of the anxiety over AI and other advanced technologies stems from the fear of the displacement of labour, as significant numbers of jobs in many industries may soon be performed by a combination of robotics, automation, and algorithms. From a business ethics perspective, this presents business leaders with a challenge: balancing the best interests of the business, the employees, and the wider community and society. It's a task made more complex by the fact advanced technology systems aren't self-sufficient; while AI can replace some jobs, for example, it creates others that often require highly specialized skills and advanced training. It's also hard to predict how AI and other advanced technologies will ultimately be used. Business needs and customer preferences and behaviours are constantly changing, driving leaders and their employees to look for increasingly innovative ways to use technology.

Embracing new, disruptive technologies brings its own risks and unknowns for business—a factor that persuades many to adopt them incrementally. Incremental adoption of advanced technologies such as AI can mean the organization experiences fewer problems, resolving those that arise through future design decisions. In fact, much of the literature suggests that incremental adoption of advanced technologies is a more responsible approach, because it permits closer supervision by human agents. Naturally, this should also include implementing these tools in ways that are entirely transparent to employees, business partners, and customers. But how do we balance these concerns with the need to adapt quickly and remain competitive?

France's recently announced national AI strategy could provide a model for businesses and support for an incremental approach to AI adoption. The French approach resembles an "ethics of care" in that it focuses on maximizing profit through stronger relationships between employers, employees, and wider society.[9] In justifying the approach, French President Emmanuel Macron stated, "AI is about disruption and dealing with impacts of disruption. For instance, this kind of disruption can destroy many jobs in some sectors and create a need to retrain people. But AI could also be one of the solutions to better train these people and help them to find new jobs, which is good for my country, and very important."[10]

Key questions

• How should business leaders approach the issues of job transformation and displacement?

• What role does government play?

• Do we need to consider how we structure a fair post-labour economy?

Removing bias and fostering inclusion

One of the biggest challenges emerging for policy-makers and developers in the adoption of AI, machine learning technology, and other advanced technologies is ensuring that these systems support inclusivity and avoid perpetuating bias. It's a significant challenge because we humans are inherently biased, and since "AI is just an extension of our existing culture,"[11] our own biases and prejudices can be unintentionally embedded into the algorithms at the heart of our AI systems.[12] For example, a company's hiring process that uses an AI system to select candidates might screen out certain sectors of the population if the information (e.g., job qualifications) for the algorithm is based on the characteristics of the traditionally dominant demographic group. This typically disproportionately affects traditionally under-represented groups such as visible minorities, women, Indigenous peoples, members of the LGBTQ+ community, and disabled people. And because this bias is already built into the system, this discrimination takes place automatically, without human input.

Some argue that humans should be held legally accountable for advanced technologies that discriminate against users or perpetuate internal or unconscious bias.[13] But this raises new risks and uncertainties: how would we identify who—or what—to hold accountable? It seems clear we need to find better ways to protect against unconscious bias or discrimination within our technologies. Yet how do we do that, when those technologies so readily reflect their designers' own conscious or unconscious biases?

A number of potential solutions have already been put forward. Businesses and technology developers can strive to create systems that are deliberately inclusive; for example, they could design hiring systems to ignore demographic information about job candidates.[14] Establishing inclusive design practices could help developers better understand and address potential elements in their products, systems, and algorithms that could inadvertently and unintentionally exclude certain individuals.[15] This often begins with development teams themselves being more diverse and better reflecting the end user.[16]

Organizations could also focus on "emotional" AI systems. Emotional AI systems help to make more inclusive decisions, since they "exhibit user awareness by confirming and, as necessary, correcting understanding of the user's intent, and by recognizing and adjusting to the people, places and events that are most important to users."[17] Governments could also step in with regulation; for example, requiring inclusivity to be incorporated into AI systems.

Key questions

- How can developers anticipate and mitigate the potentially harmful outcomes of their technology?

- How can businesses and developers ensure that new technologies don't perpetuate bias?

- How can AI be used to promote inclusivity?

Tackling accessibility and inequality

New technologies are often touted as innovations that will change the way we live, work, and do business. In reality, the benefits of new technologies are rarely shared equitably throughout society. In fact, technological innovation often leaves marginalized or disenfranchised groups behind. Today, advanced technologies disproportionately benefit privileged groups in society. Wealthier people generally have greater access to technology and therefore are able to benefit from it more than those who are less wealthy.[18] How do we counteract technology-driven inequalities and ensure those on the margins aren't left even further behind?

Ensuring technology is available to everyone, particularly those who are hardest to reach, is an important equity consideration for government. Funding for educational and social programs that provide students—both children and adults—with the tools to learn about the technology available to them is one option.[19] One example of such a program is the Accessible Technology Program, aimed at Canadians with disabilities.[20]

But business also plays a crucial role in facilitating greater access to technology. It's already been a longstanding practice for many manufacturers to offer price discounts to students, teachers, and schools for tablets and computers. More recently, technology companies have launched initiatives to increase internet access for people living in regions of the world where access to high-speed internet and data is poor (e.g., X's Project Loon and Amazon's Android browser for emerging markets).

It's vital that government and business also understand how to mitigate technology-driven inequality in the labour market—and inequality of opportunity. These challenges are closely related to the earlier discussion about worker displacement and transitions. "[M]ost of the gains from productivity growth may not be widely shared; rather, they will likely benefit capital owners and a few highly skilled workers with strong cognitive, adaptive and creative skills, who are the best equipped to work with the machines", according to the United Nations Conference on Trade and Development.[21] Recent work by the Organisation for Economic Cooperation and Development has already found evidence of a growing divide between the productivity of firms at the frontier and those lagging behind, and linkages to increasing inequality of income and opportunity.[22] Job training programs that equip displaced workers with the skills necessary for the future of work will help address the risk of growing inequality, but bolder collective action by government, academia, and business will be needed. This includes a shift in mindset from a traditional linear path (i.e., study, work, retire) to one that is non-linear and embraces life-long learning.[23]

Key questions

- Can technology be an equalizer?

- What role do governments have in ensuring new technologies are deployed in ways that benefit all citizens equitably, not only a privileged segment of the population?

- In the context of talent displacement, how far does an employer's responsibility to its employees go?

The need for collaborative governance

Good, effective governance of emerging technologies requires all relevant stakeholders—industry, consumers, businesses, governments, academia, and society—to work together in a collaborative, decentralized way.

Governments can't do it alone, nor should they. Tackling the challenge demands better coordination of ideas, resources, information, and planning. Technology experts must work alongside business and government leaders to make better, more ethical decisions.

Multi-stakeholder approaches will not only achieve the greatest economic and technological benefits, but also ensure that the broadest range social benefits are possible. Looking ahead, collaboration may in fact become a necessity as regulatory frameworks and actors become global to align with the global reach of technology, data, and information.[24] An example of collaboration is the European Union's Responsible Research and Innovation process, in which all stakeholders, including business, are involved in the decision-making processes of technologies, specifically in terms of social issues but also in terms of ethics and good governance.[25] Technology experts can anticipate potential challenges; businesses can interpret the impacts on customers or clients; and governments can address potential problems affecting society more widely.

Fostering more dialogue among technology stakeholders can also allow businesses to play a key role in helping governments as they develop laws and standards that increase the reliability of emerging technologies.[26] Academics, researchers, and technology experts can be encouraged to make their findings more accessible to business and government leaders as well as the broader public, so that informed decisions can be made and policies and guidelines established. Frank, candid discourse about new technologies, for example, could lead to new rules and guidance concerning matters of privacy, transparency, inclusivity, accessibility, inequality, and more, mitigating these issues before they become problems.

Key questions

- How should business and government collaborate to ensure good governance?

- What is the role of policy-makers and regulators in advancing technology and business?

- How should business leaders engage with policy-makers and regulators about technology issues?

- What role could Canada play in driving forward these conversations?

Conclusion

The ethics of business and technology are difficult matters to grapple with, and the relentless pace of disruptive technological change adds to the challenge for business and government leaders alike. Government finds itself reacting because of the nature of policy- and lawmaking. Business, on the other hand, can move much more quickly—and this means there is an important opportunity for business to demonstrate leadership in ensuring new technologies operate in an ethically sound manner. But what does it really mean to develop technologies that achieve "good" outcomes?

This paper is intended to spark an important dialogue at the intersection of business ethics, policy-making, and technology. Business leaders, government leaders, academics, developers, and technology experts need to come together to discuss how technology is changing our society in both positive and negative ways.

There are, of course, many more questions to ask as we try to form an understanding of these complex, ever-changing issues. We hope the questions posed—and the ideas put forth in this discussion paper— contribute to your own dialogue with your leadership teams, coworkers, and colleagues.

Endnotes

1   Herschel, Richard and Virginia M. Miori. "Ethics and Big Data." *Technology in Society* 49 (2017): 34.

2   Stapleton, Larry. Ethical Decision Making in Technology Development: A Case Study of Participation in a Large-Scale Information Systems Development Project." AI & SOCIETY, 22, no. 3 (2008) 406.

3   Herschel and Miori. "Ethics and Big Data." 31.

4   Obar, Jonathan A., and Anne Oeldorf-Hirsch. "The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network, August 24, 2016. https://papers.ssrn.com/abstract=2757465.

5   Kerr, Ian R., and Marcus Bornfreund. "Buddy Bots: How Turings Fast Friends Are Undermining Consumer Privacy." Presence: Teleoperators and Virtual Environments 14, no. 6 (2005): 653.

6   Smith, Brad, and Harry Schum. "The Future Computed: Artificial Intelligence and Its Role in Society." 2018. https://blogs.microsoft.com/uploads/2018/02/The-Future-Computed_2.8.18.pdf, 68.

7   Pieters, Wolter. "Explanation and Trust: What to Tell the User in Security and AI?" Ethics and Information Technology 13, no. 1 (2010).

8   Ibid, 55.

9   Thompson, Nicholas. "Emmanuel Macron Talks to Wired About France's AI Strategy". *Wired*, March 2018. < https://www.wired.com/story/emmanuel-macron-talks-to-wired-about-frances-ai-strategy/>.

10 Ibid.

11 Hutson, Matthew, et al. "Even Artificial Intelligence Can Acquire Biases against Race and Gender." American Association for the Advancement of Science, 26 July 2017, www.sciencemag.org/news/2017/04/even-artificial-intelligence-can-acquire-biases-against-race-and-gender.

12 Etzioni, Amitai, and Oren Etzioni. "Incorporating Ethics into Artificial Intelligence." *The Journal of Ethics* 21, no. 4 (2017): 403.

13 Etzioni and Etzioni. "Incorporating Ethics into Artificial Intelligence." 413.

14 Knight, Rebecca. "7 Practical Ways to Reduce Bias in Your Hiring Process." *Harvard Business Review.* July 05, 2017. https://hbr.org/2017/06/7-practical-ways-to-reduce-bias-in-your-hiring-process.

15 Smith, Brad, and Harry Schum. "The Future Computed", 69.

16 Deloitte, "Outcomes over Optics: Building Inclusive Organizations." November 2017. https://www.canada175.ca/sites/default/files/download/files/inclusion_aoda_en_0.pdf.

17 Smith, Brad, and Harry Schum. "The Future Computed," 70.

18 Toyama, Kentaro. "Will Job-stealing Robots Cause a Human Revolution?" November 2017. https://www.weforum.org/agenda/2017/11/will-ai-job-stealing-robots-lead-to-a-human-revolution.

19 Baller, Silija, Soumitra Dutta, and Bruno Lanvin. "The Global Information Technology Report: Innovating in the Digital Economy." 2016. http://www3.weforum.org/docs/GITR2016/WEF_GITR_Full_Report.pdf, 7.

20 Innovation, Science and Economic Development Canada. "Accessible Technology Program." Canada.ca. November 30, 2017. http://www.canada.ca/en/innovation-science-economic-development/programs/accessible-technology.html.

21 United Nations Conference on Trade and Development. "Information Economy Report 2017: Digitalization, Trade and Development." 2017. http://unctad.org/en/PublicationsLibrary/ier2017_en.pdf. 67.

22 Organization for Economic Cooperation and Development. " The Productivity Inclusivity Nexus" June 2016. https://www.oecd.org/global-forum-productivity/library/The-Productivity-Inclusiveness-Nexus-Preliminary.pdf.

23 Deloitte. "Bold Bets for Our Country: It's Time for Deliberate Action." Deloitte, July 2017, https://www2.deloitte.com/content/dam/Deloitte/ca/Documents/press-releases/ca-final-aoda-bold-bets-en.PDF.

24 Chakravorti, Baskar. "Why the Rest of the World Can't Free Ride on Europe's GDPR Rules" *Harvard Business Review.* April 30, 2018. https://hbr.org/2018/04/why-the-rest-of-world-cant-free-ride-on-europes-gdpr-rules.

25 Ibid, 14.

26 Smith, Brad, and Harry Schum. "The Future Computed", 64.

Acknowledgements

Deloitte.

www.deloitte.ca