


# Responsible AI in Financial Services

Enhancing Trust While  
Adopting New Capabilities



# Introduction



At TD, we believe that innovation – like banking – can enrich lives.

There is little doubt that artificial intelligence (AI) is among the most powerful new innovations in the market today. It can organize billions of data points into clear insights, create personalized experiences, anticipate needs and be used to manage risks associated with cybersecurity and fraud.

As the financial services industry embraces the power of this new technology, it is important to continue reflecting on and analyzing the risks and barriers – big or small – to its adoption and acceptance.

It is our belief that the potential benefits of AI are enormous, both for our industry and across society. When used responsibly, it can improve lives when applied to healthcare, reduce poverty when harnessed by governments, provide customers with better services, and solve problems by finding hidden insights in data.

However, for this future to be realized, the financial services industry must map the evolution and enable the adoption and use of AI in a responsible way. While individuals are eager for more relevant and highly personal experiences, our research shows that they are wary of the risks inherent in any form of

automation. As an industry, it is imperative that we continue to identify these types of concerns, raise the challenges and seek a path forward that is both effective and responsible.

Like many new and emerging technologies, AI suffers from a lack of understanding. Our research shows that businesses and consumers lack clarity on what AI is, what it can do, and where opportunity and risk exist. Demystifying AI is the first step.

We therefore initiated the work for this report. Not to solve every problem or concern, but to advance an ongoing dialogue. To put the issues on the table and ignite a conversation in an open and transparent way. To bring together leading thinkers and primary research to both enhance our understanding of the potential of AI and think through the concerns that could impede progress.

We believe that the good from AI will far outpace the risks. But only if we progress through open dialogue and move towards a more prosperous future in a responsible way.

**Michael Rhodes**, *Group Head, Innovation, Technology and Shared Services*


# About the Report

Considering the importance of banks to the economy and the lives of Canadian families, financial institutions can support an open dialogue on the potential uses of AI.

TD is leveraging its expertise in both banking and AI to advance this conversation, which we hope will help foster the responsible use of this technology in financial services.

TD recently engaged a cross-section of experts to participate in a roundtable to better understand the risks associated with AI in financial services and, in preparation, commissioned a national survey of Canadians to better understand their use of, and trust in, AI.

This report brings the insights from that roundtable together with the perspectives of those Canadians to inform key areas where the financial services industry needs to focus as we work to build out best practices for the responsible use of AI.



---

“We are in a country that is at the forefront of AI and we need to work together to map a path forward for the responsible use of AI. It’s a great time for Canadian industry to lead.”

*Tomi Poutanen, Chief AI Officer, TD  
and Co-founder of Layer 6*

# What Canadians Say

TD's survey of Canadians reveals that the majority are comfortable with AI if it means they will receive better services. However, they have concerns.



## Awareness

Canadians tell us that they believe they are interacting with AI daily, and they recognize the value that AI adds to their lives:

**72%** are comfortable with AI if it means they will receive personalized services



## Innovation

Canadians expect their banks to continue to keep pace with new technology:

**87%** say they want banks to be innovative

**78%** are willing to try new forms of technology from financial services companies, including apps and online services

**72%** say they have come to expect personalized experiences from the companies they transact with



## Concerns

Despite their awareness of AI, Canadians have concerns:

**77%** say they are concerned about the risks AI poses to society

**77%** are concerned that AI is advancing too quickly to really understand the potential risks

**68%** are concerned they don't understand the technology well enough to know the risks

**59%** believe that the lack of diversity among people working in the field of AI could lead to biases in how the technology is being developed



## Expectations

When asked for their views on factors that are important when it comes to how companies use AI, Canadians cite the following:

**70%** say customers should have control over how their data is used

**55%** say companies should be transparent about how they are using the technology

**28%** say decisions made using AI should be easy to explain and understand



# Insights from the Roundtable



TD brought together a group of experts in AI from across multiple sectors, including financial services, technology, FinTech, academic, public and not-for-profit organizations, for a roundtable to explore the unique opportunities and challenges facing the financial sector in the adoption of AI.

The roundtable participants analyzed future-state scenarios that presented instances where AI resulted in unintended consequences for customers. Throughout the conversation, the group identified the following areas of focus for financial institutions as they adopt and develop best practices on the responsible use of AI:

**Explainability**


Determine how AI experts and business leaders should approach the inherent limitations of the technology as it relates to explaining how AI models arrive at conclusions.

**Bias**

Control for bias and re-examine the concepts of transparency, fairness and accountability in an AI-first world.

**Diversity**

Promote the role that diversity and inclusion should play at every level of AI adoption, from executive leadership to the teams building AI models, to the data used to inform decisions.



Brent Barron  
**CIFAR**

Alex LaPlante  
**Borealis AI**

Jason Millar  
**University of Ottawa**

Mike Shaver  
**integrate.ai**

Alice Yang  
**TD Bank Group,  
Layer 6**

Sasha Caskey  
**Kasisto**

Maithili Mavinkurve  
**Sightline Innovation**

Tom Oddie  
**TD Bank Group**

Elissa Strome  
**CIFAR**

Richard Zemel  
**University of  
Toronto,  
Vector Institute**

Matt Fowler  
**TD Bank Group**

Charles McCarragher  
**TD Bank Group**

Tomi Poutanen  
**TD Bank Group,  
Layer 6**

Steve Sweetman  
**Microsoft**

Kathryn Hume  
**Borealis AI**

Kathleen McGinn  
**Rotman School of  
Management**

Hossein Rahnama  
**Flybits,  
Ryerson University**

Jodie Wallis  
**Accenture**

# Explainability

As roundtable participants worked through scenarios that addressed issues of bias, risk, and human judgement, a common thread emerged around the importance of communication and consensus between executives and scientists on the capabilities of AI and its shortfalls. To uphold accountability and transparency, it is critical to have open communication and a common understanding among team members of what AI is and what AI can and cannot do.

Communication barriers that may exist between executives and engineers, or between companies and customers, can be linked back to the idea of ‘explainability’ in AI — the ability to explain how an AI system arrived at a conclusion. This may also include explaining to individuals when and how decisions are being made about them using AI. Explaining an output of an AI system is not always straightforward when dealing with extremely large volumes of data and complex methods of processing that data. The roundtable experts broke down the challenge of explainability into three key components:

## Identifying what needs to be explained

The need to explain decisions made through AI can vary in importance and companies should prioritize efforts on areas that have the potential to negatively influence or impact people’s lives.

## Expecting the unexpected

Machine learning models can change over time as they ‘learn’ or have new data inputted, and this may happen in ways that people may not anticipate. When addressing explainability, companies should implement processes and standards that evolve alongside their models and continuously test for inconsistencies.

## Educating and reaching consensus

Technologists, government and business leaders need to come to a clear and agreed upon understanding of the technical capabilities and limitations of AI models so that realistic expectations can be set around explainability, transparency and accountability.



**of Canadians are concerned  
they don't understand AI well  
enough to know the risks**



**“The need to explain decisions made by AI is important** when the decisions affect someone’s life. For example, you need to be prepared to explain how you made a hiring or firing decision, whether you extended credit or were provided an insurance policy. Other decisions may not have the same requirements for explainability. Companies need to keep that distinction in mind.”

*Jodie Wallis, Managing Director, Artificial Intelligence, Accenture*

**“There’s a lot of Hollywood-type images of the ‘robots are coming.’** But machine learning is already part of our everyday lives. People don’t have to understand all the technical details, but they do need to understand the risks and the benefits. AI is going to transform – for the better – our medical system, our transportation system, help us to mitigate climate change, and so much more.”

*Elissa Strome, AVP Research and Executive Director, Pan-Canadian AI Strategy, CIFAR*

**“Accountability, when it comes to eliminating discrimination from AI decisions,** isn’t just one person or one decision point. It’s all over the organization and must reflect your company values and the values of society.”

*Sasha Caskey, CTO and Co-founder, Kasisto*



# Bias


Most people would agree that AI systems and solutions must be fair and unbiased so that they do not create an unintended disproportionate or negative effect on some groups of users relative to others. The question of how to address bias was a key theme at the roundtable:

## Bias can have multiple meanings

Bias is a term that can mean different things in different contexts and to different people. Generally, the concern around biased decision-making in AI models stems from the concept of human bias, which can lead to unfair treatment or discrimination, and must be avoided or mitigated. However, statistical bias, for example, has a different meaning and can be useful in an AI model. When companies examine their processes to guard against bias in AI, they must consider perspectives from varying types of expertise.

## Data makes a difference

The data used in AI models is a component affecting whether a model might produce a biased outcome. For example, a model could be accurate across the Canadian population but may produce different outcomes if applied only to a specific sub-population. Technologists must consider this dynamic, as well as carefully consider whether historical discrimination might be present in the data used.



**70%**  
of Canadians list  
having control over  
how their data is used  
as a top concern when  
it comes to AI

## Bias may show up in unexpected places

The roundtable participants noted that when one characteristic – such as gender, age or ethnicity – is removed from data to eliminate biased outcomes, machine learning models will often create proxies for that same characteristic.

**“We need to acknowledge that bias will exist,** and will continue to exist, and we need the right processes and procedures in place to address it on an ongoing basis. If we continue to talk about bias as this amorphous thing without really defining it, then I think we’re missing the opportunity to develop better models and make better decisions.”

*Jodie Wallis, Managing Director, Artificial Intelligence, Accenture*

**“Balance is key when designing AI systems.** While removing data to protect against bias or to uphold privacy standards may be necessary in some cases, a lack of data can cause problems in the AI model. Financial services need an approach that takes this challenge into consideration.”

*Richard Zemel, Professor in Computer Science, University of Toronto and Research Director, Vector Institute for Artificial Intelligence*

**“Each one of us bring our own biases** and we need to understand the impact of that bias in the data which we use in our systems. The most effective way to truly address the impact of discriminatory bias in datasets today is through transparency, human oversight and accountability.”

*Steve Sweetman, Director, Office of Responsible AI, Microsoft*

## From the Experts



# Diversity

Roundtable participants reflected on one of the biggest challenges in the adoption and implementation of AI: diversity and inclusion. The experts agreed there is a need for the inclusion of diverse perspectives in the development of AI systems so that the models reflect the audience being served. The following areas were identified as critical for organizations to consider:

## Building diverse AI teams

From the engineers to the executives, teams should represent the customers they serve.

## The need for representative data sets

Diversity and representation matters when choosing data sets to train AI models, and the people being served by AI systems must be represented in the data.

## Canada can lead

With a multicultural society and workforce, Canada has an opportunity to play a leading role on a global stage when it comes to fostering diversity and inclusion in the AI sector.



# 59%

of Canadians believe that the lack of diversity among people working in the field of AI could lead to biases in the technology being developed



**“One of the big challenges right now around the adoption of AI in Canada** and around the world is about equity, diversity and inclusion. Not only in researching and implementing, but also benefitting from AI systems. If we want AI to have a positive social impact for all of humanity, we need to address these disparities.”

*Elissa Strome, AVP Research and Executive Director, Pan-Canadian AI Strategy, CIFAR*

**“Recognizing that there’s an increasing population of newcomers to Canada,** we need to consider the downstream effects on the economy if their data is left out of the financial systems informed by AI.”

*Kathryn Hume, Director, Product and Business Development, Borealis AI*

**“Canada attracts top AI talent from around the world,** and the field of AI here has really thrived because of this diverse talent pool. In fact, the majority of my team comes from countries outside of Canada. It’s critical that we continue to bring diverse perspectives to the table that challenge our thinking.”

*Tomi Poutanen, Chief AI Officer, TD and Co-founder of Layer 6*

## From the Experts



# A Look to the Future



# Final word from Michael Rhodes

The trust that our customers place in us is central to our innovation philosophy, no matter which set of technologies we're exploring. As the adoption of AI develops and continues across the financial service industry, we believe this is a critical time to advance an industry-wide discussion that moves beyond principles to create world-class services for Canadians in a responsible way, while unlocking the full potential of AI. We see this report as an important step in that direction.

At TD, we believe we have a duty to work with industry experts in the field of AI to understand how we can protect the interests of our customers as we adopt new and innovative ways to meet their needs.

The roundtable, together with the survey, revealed that while we have been taking positive foundational steps, the financial services industry has some work to do. Our work must focus on how we can establish

truly diverse teams and how we can advance the collective process to explain the potential and the limits of what is possible in an AI-first world and to dig deeper into the questions remaining about the role of bias.

We believe that AI can be an important part of our future efforts to support our customers, colleagues and communities, and we must work to develop this technology so that all Canadians benefit. We are excited by what promises to be one of the world's most influential technologies. And we look forward to the continued discussions across our bank and the financial services industry as we come together to chart a responsible and prosperous path forward.

A handwritten signature in black ink, appearing to read "Michael Rhodes". The signature is fluid and cursive, with a long horizontal stroke extending to the right.

# Appendix:

## Responsible AI Roundtable Participants

### Brent Barron

Director, Public Policy  
CIFAR

Brent Barron is responsible for engaging the policy community around cutting-edge science. He oversees CIFAR's AI & Society program, examining the social, ethical, legal, and economic effects of AI.

### Alex LaPlante

Head of Risk and Commercial Banking  
Borealis AI

Dr. Alex LaPlante leads Borealis AI's efforts to develop and implement AI solutions to improve efficiency and flexibility across the risk and commercial banking lines of business.

### Sasha Caskey

CTO and Co-Founder  
Kasisto

Sasha Caskey has spent his career teaching computers to communicate with humans. He has driven the development of the KAI platform since it spun out from SRI International, the world-renowned research institute where Siri was invented.

### Matt Fowler

Vice President, Enterprise Machine Learning  
TD Bank Group

Matt Fowler is responsible for developing and leading a team within Enterprise Data and Analytics (ED&A) to identify, scope, size and prioritize machine learning opportunities across the bank's lines of business at TD.

### Kathryn Hume

Director, Product and Business Development  
Borealis AI

Kathryn Hume is passionate about building companies and products that unlock the commercial value of emerging technologies. She has held leadership positions at integrate.ai and Fast Forward Labs (Cloudera), where she helped companies apply machine learning to increase revenue and operational efficiency.

### Maithili Mavinkurve

Founder and COO  
Sightline Innovation

Maithili Mavinkurve is one of the first female founders of an AI company in Canada. She has also represented Canada at the G7 ministerial meetings on AI and the future of work. She was recently acknowledged as one of the "30 Most Influential Women in AI in Canada."

### Charles McCarragher

Vice President, Legal  
TD Bank Group

Charles McCarragher leads the Global Contracts, Global Real Estate, Channels, Payments and Marketing Legal teams supporting TD Bank's legal needs in these areas. He has a special expertise in working with innovative companies of all sizes across various product and service verticals.

### Kathleen McGinn

Instructor and DBA Candidate,  
Rotman School of Management

Kathleen McGinn works and studies at the intersection of Ethical Stewardship, Strategy and Economics and is currently pursuing a Doctoral Degree in Business Administration jointly at The Henley Business School in London, England and the University of Toronto.

### Jason Millar

Canada Research Chair in the Ethical Engineering of Robotics and Artificial Intelligence, University of Ottawa's School of Electrical Engineering and Computer Science

Jason Millar researches the ethical engineering of robotics and AI, with a focus on developing tools and methodologies engineers can use to integrate ethical thinking into their daily engineering workflow. Jason recently co-authored a discussion paper as one of Canada's contributions to the 2018 G7 meeting on AI in Montreal.

### Tom Oddie

Vice President, Product Analytics, Canadian Personal Banking  
TD Bank Group

Tom Oddie leads the analytics function for personal banking, helping to realize the team's vision to make analytics a competitive advantage for the personal bank. Tom is a career data scientist, with a post-graduate degree in Statistics from the University of Warwick (UK).

### Tomi Poutanen

Co-founder, Layer 6  
Chief AI Officer, TD Bank Group

Tomi Poutanen is TD's Chief AI Officer and Co-founder of Layer 6. He is also a Co-founder of the Vector Institute for Artificial Intelligence, a leading academic research institute for deep learning, and a founding fellow at the Creative Destructive Lab, the world's largest AI venture accelerator, located at the University of Toronto.

### Hossein Rahnama

Founder and CEO  
Flybits

Hossein Rahnama is a computer scientist who lives in the paradox of academia and entrepreneurship. Beyond his role at Flybits Inc., Hossein is a visiting professor at the MIT Media Lab and professor at Ryerson University.

### Mike Shaver

CTO  
integrate.ai

Mike Shaver oversees engineering, design, product, and machine learning research in pursuit of new ways to help businesses become more customer-centric. Mike was previously a partner at Real Ventures, led Facebook's transformation to a mobile-first company and managed Firefox development as VP Engineering at Mozilla.



**Elissa Strome**

AVP Research and Executive Director, Pan-Canadian AI Strategy, CIFAR

Elissa Strome works with our three national AI Institutes and researchers across the country to advance Canada's national AI research and innovation strategy. She holds a PhD in Neuroscience from UBC.

**Steve Sweetman**

Director, Office of Responsible AI  
Microsoft

Steve leads the enablement strategy for Responsible AI at Microsoft, putting Responsible AI principles into practice, both within Microsoft, and through tools and practices that empower customers and partners to do the same.

**Jodie Wallis**

Managing Director, Artificial Intelligence  
Accenture

Jodie Wallis is responsible for expanding Accenture's business presence in the AI ecosystem, attracting top talent and bringing innovation and global AI expertise to Canadian clients across all industries. She leads Accenture's collaboration with business partners, government and academia and oversees its investments in the Canadian AI ecosystem.

**Alice Yang**

Chief of Staff, Layer 6  
Senior Manager, TD Bank Group

Alice Yang oversees the planning and execution of machine learning projects and manages financial and business operations at Layer 6. She also has consulting experience working with tech start-ups and companies including Google and SAP.

**Richard Zemel**

Professor in Computer Science, University of Toronto  
Research Director, Vector Institute for Artificial Intelligence

Richard Zemel is a Professor of Computer Science at the University of Toronto and Research Director at the Vector Institute. He is an Industrial Research Chair in Machine Learning, and also the co-founder of SmartFinance, a financial technology start-up specializing in data enrichment and natural language processing.



