CIVIC
IMPACT / COALITION

Capital with conviction.
Impact with returns.

BUILDING THE NEXT GENERATION OF IMPACT INVESTING IN CANADA

Building AI-Native Impact Investing: AI Agents, Governance & the Future of Decision-Making

A CIVIC Impact Investing Circle on how AI agents are reshaping the architecture of investing — and why human judgment and governance become more valuable, not less.

ABOUT THIS REPORT

Contributors & organizations

EDITOR

Lyne Landry — Co-founder & Managing Partner, AgeTech Capital; Co-founder & Chair, CIVIC

SUBJECT EXPERT

Nicholas Andreou, Head of AI, Better Society Capital

IN COLLABORATION WITH

Better Society Capital, ImpactVC and Circle contributors

About CIVIC

CIVIC is a Canadian impact investing community and platform dedicated to advancing the next generation of impact capital, collaboration, and thought leadership. Through curated investor circles, strategic convenings, research initiatives, and ecosystem-building activities, CIVIC brings together experienced investors, fund managers, institutions, family offices, founders, and ecosystem leaders to explore the evolving challenges and opportunities shaping impact investing.

CIVIC's work focuses on strengthening the sophistication, rigor, and long-term effectiveness of impact investing in Canada and globally. Its circles and initiatives address emerging topics such as AI governance, systems change, impact risk assessment, climate innovation, inclusive finance, and the institutionalization of impact management practices.

By creating spaces for candid dialogue, peer learning, and cross-sector collaboration, CIVIC aims to help build a more resilient, accountable, and innovation-driven impact investing ecosystem.

About Better Society Capital & ImpactVC

ImpactVC is a community and knowledge platform for venture capital leaders working at the intersection of impact and investment. It equips fund managers and investors with frontier insight, tools, and trusted relationships to raise and deploy capital more effectively. Through events, resources, and peer learning, ImpactVC fosters high-quality dialogue between VCs and LPs on impact practice. It brings together those who believe that some of the largest market opportunities of the coming decades lie in tackling complex social, environmental, and economic challenges.

Better Society Capital (BSC) is the UK's leading impact investor, dedicated to growing the social impact investment market. In venture, BSC channels capital into VC funds and scalable impact ventures — backing innovative technology solutions addressing pressing social challenges such as financial inclusion, health, and education. Through its venture work, BSC seeks to demonstrate that backing mission-driven founders can deliver both strong financial returns and meaningful social change.

CONTENTS

In this report

FOREWORD

Market standards are not built in isolation. They emerge when practitioners have the trust to speak candidly, challenge assumptions, and share what is and is not working.

Impact Circles creates that space. By bringing together a small group of 12 to 15 peers, the Circles foster the trust, openness, and active participation needed to tackle the questions shaping impact investing in Canada and build greater credibility around the ability to achieve market-rate returns and impact at scale.

This report carry the insights beyond the room and turning practitioner experience into shared knowledge that can strengthen practices, inform market standards, and advance the entire ecosystem.

This work would not be possible without the generous support of Realize Capital, Fonds de finance sociale CAP Finance, and Blue HF Legal, and the leadership of CIVIC, Better Society Capital and ImpactVC. Their commitment enables the ecosystem to create the trusted spaces it needs to learn and evolve together.

EXECUTIVE SUMMARY

Artificial intelligence is rapidly reshaping the practice of investing itself.

What began as experimentation with workflow automation, research support, and productivity enhancement is evolving into something far more profound: the emergence of AI-native investment organizations. Across venture capital, private equity, and financial services, firms are beginning to redesign the architecture of decision-making around large language models ("LLMs"), AI agents, retrieval systems, and continuously learning knowledge infrastructures.

This CIVIC Impact Investing Circle was led by Nicholas Andreou, Head of AI from Better Society Capital, who shared the organization's experience building an internal AI-powered impact assessment agent and integrating AI systems into investment workflows. The session offered participants a rare operational perspective into how established impact investors are beginning to move beyond generic AI experimentation toward the construction of proprietary, governance-driven AI infrastructures capable of supporting investment and impact assessment processes at scale.

Throughout the discussion, participants explored how AI is being integrated across nearly every stage of the investment lifecycle — from sourcing and market mapping to due diligence, memo generation, portfolio monitoring, reporting, and internal knowledge management. While the operational gains appear extraordinary, the discussion quickly evolved beyond efficiency. Participants repeatedly returned to a deeper realization: AI is not simply becoming another tool inside investment firms. It is beginning to influence how investment judgment itself is formed.

This introduces both transformative opportunity and significant new systemic risk.

Participants discussed how AI systems may dramatically improve analytical speed, institutional memory, workflow scalability, and pattern recognition. At the same time, the Circle surfaced growing concerns around hallucinations, model reliability, consensus investing, explainability, governance, and the erosion of differentiated thinking.

As one participant observed:

> "Once everyone uses the same LLMs for investing, there's a risk that everyone converges toward the same conclusions — exactly when investing requires differentiated thinking."

The discussion strongly aligned with broader shifts now emerging across institutional investing: the rise of AI-native operating models, proprietary knowledge systems, human-AI collaborative workflows, and governance architectures designed specifically for AI-assisted decision-making.

A central insight emerged repeatedly throughout the Circle: the firms that win in an AI-native world may not be those that automate the most, but those that most effectively combine machine intelligence with human judgment, systems thinking, governance, and trust.

The group ultimately concluded that AI agents may fundamentally reshape investment workflows over the next decade, but that human discernment will likely become even more valuable in areas requiring contextual reasoning, ethical judgment, contrarian thinking, and governance oversight. As AI systems become embedded into capital allocation itself, institutional-quality AI governance may increasingly become part of fiduciary responsibility.

01 AI agents and the future operating model of impact investment firms

Historically, investment firms scaled through analyst headcount, proprietary networks of experts, specialized investors, sector specialization, and accumulated institutional knowledge. Much of the investment process depended on labor-intensive research, fragmented information systems, and tacit knowledge residing inside individual partners or teams.

Participants described how modern LLMs and AI agents are now capable of synthesizing thousands of pages of documents, mapping markets, comparing competitors, generating first-draft investment memos, surfacing anomalies, retrieving historical decisions, and organizing institutional knowledge almost instantly. Several participants discussed how their organizations are already experimenting with internal AI academies, proprietary retrieval systems, AI copilots, and knowledge architectures connected to internal investment databases. The implications are profound. Investment firms may increasingly resemble continuously learning intelligence systems rather than purely human-driven analytical organizations.

However, the Circle argued that uniquely human capabilities — including trust-building, founder assessment, systems intuition, narrative formation, ethical reasoning, and governance oversight — may become even more important as baseline analytical capabilities become commoditized.

This surfaced one of the defining strategic questions of the session:

> "If information itself becomes increasingly democratized through AI, what becomes the true source of investment edge?"

Participants suggested that future differentiation may increasingly depend on proprietary workflows, organizational learning systems, differentiated worldviews, and the ability to structure and govern institutional knowledge more intelligently than competitors.

At the same time, the Circle strongly emphasized that AI amplifies both strengths and weaknesses. Participants stressed that building reliable AI systems requires extensive iteration, evaluation, testing, prompt engineering, governance discipline, and continuous validation. The discussion reinforced the idea that access to models themselves is unlikely to become the primary competitive moat, particularly as foundational models become increasingly commoditized.

Instead, future advantages may depend on system architecture, proprietary data, workflow integration, organizational learning, and governance quality.

02 Reliability, hallucinations & the governance challenge

One of the strongest discussions during the Circle centered on reliability. Participants repeatedly questioned how investors can trust AI-generated analysis when LLMs are capable of hallucinating, fabricating, or misinterpreting information.

The Better Society Capital AI specialist emphasized that building a reliable institutional-grade AI investment system is far more complex than simply connecting a large language model to public information. The quality of the outputs ultimately depends on the quality of the underlying data architecture, quality of prompt, governance systems, validation processes, and human oversight structures surrounding the model itself.

One of the first challenges discussed was data quality. Public information alone is often insufficient for meaningful investment or impact assessment, particularly in sectors such as healthcare, climate, caregiving, education, or financial inclusion where many of the most important indicators are not publicly disclosed.

The Circle discussed how investment firms may increasingly require proprietary internal datasets built from founder interviews, proprietary dataset analysis, diligence conversations, investment memos, expert reviews, historical portfolio outcomes, and internally validated impact metrics.

Several participants emphasized that AI systems also require carefully curated reference frameworks. Without retrieval controls and structured sourcing protocols, LLMs may pull information from unreliable or low-quality internet sources. As one participant joked, investors do not want their impact assessments "sourced from Reddit."

Instead, participants stressed the importance of grounding AI systems within carefully selected institutional reference materials, peer-reviewed research, evidence-based healthcare standards, regulatory guidance, and internally approved benchmark datasets.

The Circle repeatedly emphasized that robust frameworks remain essential even in an AI-native world. AI may accelerate assessment processes, but it does not replace the need for clearly defined methodologies for evaluating impact quality, additionality, stakeholder outcomes, governance maturity, or systemic risk. Examples of validated impact frameworks include the IMP Five Dimensions of Impact from Impact Frontiers.

Participants also stressed that institutional AI systems require extensive testing before deployment. Several contributors described iterative validation processes where teams continuously compare AI-generated conclusions against expert judgment, refine prompts and workflows, and stress-test outputs across multiple scenarios.

One participant remarked that AI systems "cannot simply be left on overdrive."

Instead, firms may increasingly require continuous random testing procedures designed to verify output quality, detect model drift, identify degradation over time, and ensure ongoing reliability.

The Circle also explored systems capable not only of generating conclusions, but of evaluating the quality of their own assertions. Participants discussed models trained to classify statements or investment claims along a spectrum such as weak, moderate, strong, or evidence-backed, while simultaneously providing rationale for those rankings.

Ultimately, the discussion reinforced a critical point: AI should support judgment — not replace it.

Even highly sophisticated AI systems remain dependent on governance structures, expert oversight, contextual understanding, and real-world domain expertise. Participants repeatedly stressed that there is still "nothing like real-life experience from experts in a domain."

This was viewed as especially true in sectors such as healthcare, aging, climate, education, caregiving, and public systems, where tacit knowledge, lived experience, and contextual reasoning often shape the most important investment judgments.

The Circle concluded that the future of AI-native investing is unlikely to involve fully autonomous decision-making systems. Instead, the most credible model may be one where AI handles scale, synthesis, retrieval, and pattern recognition, while human experts provide contextual judgment, ethical reasoning, systems understanding, and final accountability.

In this emerging model, governance becomes the bridge between machine intelligence and responsible capital allocation.

03 Human judgment in an AI-native world

One of the most provocative themes of the Circle was the possibility that AI could unintentionally increase consensus thinking across the investment industry. Participants noted that if investors increasingly rely on similar models trained on similar public datasets, investment decisions may begin to converge.

This creates a paradox. The purpose of investing is often to identify non-obvious opportunities, asymmetric insights, and differentiated theses before markets recognize them. Yet AI systems may naturally optimize toward statistical consensus, pattern reinforcement, and dominant narratives.

As one participant observed:

> "Once everyone uses LLMs for investing, it will lead to general consensus while what you want as an investor is to be contrarian."

The Circle explored whether future investment differentiation may instead depend on proprietary internal data, differentiated worldview formation, systems thinking, behavioral insight, access to a network of experts, and human interpretation.

Several participants argued that AI may raise the baseline quality of investing overall while simultaneously increasing the importance of judgment, creativity, courage, and intellectual independence. This may become particularly important in impact investing, where many of the most important questions are systemic, ambiguous, values-based, and difficult to reduce into purely quantitative frameworks. Despite the enthusiasm around AI capabilities, the Circle strongly rejected the idea that investing will become fully automated.

This is particularly true in areas involving founder evaluation, trust, governance, systems change, organizational resilience, ethical ambiguity, and long-term societal impact. AI can synthesize information. But participants questioned whether it can truly understand human motivation, detect integrity, navigate power dynamics, assess moral trade-offs, or anticipate second-order societal consequences.

The Circle highlighted that impact investing often requires navigating incomplete information, contested values, stakeholder tensions, and uncertain futures. These are not purely computational problems. Participants increasingly framed the future not as human versus AI, but rather: human judgment amplified by AI systems.

The implication is important: the firms that succeed may not be those that automate the most, but those that best combine machine intelligence, governance, systems thinking, and human discernment.

04 AI governance as the next frontier of impact investing

The Circle strongly reinforced a theme already emerging across the first CIVIC report: AI governance is becoming part of impact governance.

Historically, many impact investors focused primarily on outputs, ESG metrics, and impact reporting. However, AI-enabled systems create entirely new categories of risk involving opacity, bias, explainability, accountability, worker displacement, surveillance, and automated decision-making.

These concerns are especially significant in sectors involving healthcare, caregiving, education, workforce management, insurance, and financial services.

CONCLUSION

The Circle demonstrated that AI is not simply improving investment workflows: it is beginning to reshape the architecture of investing itself.

Investment firms are evolving from human-centered analytical organizations toward AI-native intelligence systems capable of continuously compounding institutional knowledge and accelerating decision-making.

This transformation may unlock extraordinary efficiency, scalability, institutional memory, and analytical capability. But it also introduces profound new challenges around governance, reliability, explainability, consensus formation, and societal accountability.

Ultimately, the Circle reinforced a defining insight for the future of AI impact investing: it requires dedicated AI governance frameworks alongside traditional financial and operational diligence. AI governance is no longer simply a technical issue. It is becoming a fiduciary issue.

GET INVOLVED

Help build the next generation of impact investing

CIVIC's Circles, convenings, and research are built by and for their members. Join a Circle, partner on the next report, or bring your firm into the conversation.

WEB

impactvc.ca

CONTACT

info@impact-coalition.ca

SERIES

Report 2 of 3

All our reports are available on the web at impactvc.ca

CIVIC
IMPACT / COALITION

© 2025–2026 CIVIC — Canadian Impact Venture Investors Coalition.