

Equities

Corporate Engagement Spotlight: Ethical AI Practices

🕒 April 22, 2021

👤 Katherine Kroll; Lauren Cahalan

☰ TOGGLE SIDEBAR

Right now, as you read this, there are computers and network servers all over the world pondering one topic: you.

They are improving your life invisibly—streamlining traffic on your commute, helping drive down your utility costs, curating playlists for you on your favorite streaming app, possibly improving the health care diagnoses you receive. They are also watching your face, the way you walk, the color of your skin; they are monitoring your financial habits, gauging if you are a public safety threat, reviewing your social media activity, and sending you content that might shift your political stances.

This reality is the basis for one of the most powerful technological revolutions in history: artificial intelligence. The potential of AI is seemingly limitless; [AI solutions could contribute \\$15 trillion to the global economy in 2030](#)—more than the current economic output of China and India combined. AI is poised to transform the way we work, travel, run our households, power our cities and organize our society.

We are enthusiastic investors in AI—specifically in a number of leading developers of essential AI technology, as well as users of AI applications (in a separate upcoming publication we plan to discuss in greater detail the opportunities we see in this space). As with all of our investments, we seek to combine fundamental, ESG and investigative research to create a complete understanding of our holdings—the return opportunities they offer us, as well as the risks that may cause them to stumble. When we see a chance to offer feedback to these companies about the risks they face, we consider it our duty to engage with them and encourage them to consider additional ideas to improve their results and avoid harm to their stakeholders.

A number of salient risks face companies in the AI arena. Information is power, and it can be used for good or ill. Companies may misuse information—intentionally or accidentally—in many different ways, and often in recent history, the most vulnerable segments of society have borne an outsized proportion of the consequences. These risks are a potential threat not just to makers and users of these technologies, but society at large. As sustainable investors, we engage current and potential portfolio companies in discussions about these risks, to help ensure that they are managing and mitigating them in a comprehensive manner.

Artificial Intelligence: Definitions

Artificial intelligence (AI) is a broad term encompassing technologies that simulate human intelligence in machines. This simulated intelligence can manifest in several ways including, learning, problem-solving or approximation of human behavior.

AI offers immense promise as a tool to greatly improve many services and solutions used by corporate customers and benefiting society at large.

Several key concepts within AI include:

- **Machine learning (ML):** A subset of AI that reprograms itself (improving accuracy over time) as it learns from more data.
- **Deep learning:** A subset of ML where algorithms inspired by the human brain (neural networks) learn from large amounts of data, without human intervention.

Ethical AI is a term used throughout this discussion, loosely defined as development and delivery of AI products and services that empower employees, produce fair and just outcomes for stakeholders and generate other positive outcomes for society.

Assessing the Risks

AI is a broad term, and it brings with it a broad pool of risks. By focusing on transparency, protection and bias risks, our ESG research team seeks to assess a company's ability to proactively respond to potential harm before it happens. We also engage with company management teams to encourage best practices for mitigating risks, and to monitor implementation of those practices.

Transparency is needed to understand what is happening in AI models. These complex systems can be a black box for stakeholders seeking to understand them, unless companies provide detailed information about the entire process of collecting, creating, using, storing and sharing information. While such disclosures are essential, companies are seeing that transparency can generate new risks—releasing additional information about systems can make them more vulnerable to cyberattacks, and/or create potential grounds for lawsuits or regulatory action.

Finding the right balance of transparency is not a problem that can be solved overnight, but we believe that companies, investors and regulators can work together to drive gradual improvements. Beyond the developers and users of AI that volunteer an open window into their operations, we need outside actors, such as regulators and investors, to promote transparency more broadly.

Data Privacy and Protection: Data is the beating heart of any AI system, and modern society's track record on protecting consumer data leaves much to be desired. The powerful ways in which AI can use private data makes privacy and protection of that data even more important.

AI technology is creating an entirely new dimension of data-related risks. Complexity makes it almost impossible to offer viable choices to consumers about data collection, sharing and usage; it is inherently impractical for people to specifically choose to share or not share information with thousands of different online entities, or to specifically permit or deny use of that information for hundreds of different AI applications.

Another issue is fairness—many advocates argue that regulations should reflect the value of the data shared by consumers and used in AI systems and algorithms. Calculating that value is extremely difficult; some data points (like your DNA) are more valuable than others (like your email address), and the value of data changes depending on who is collecting it and how they are using it.


Finally, advanced or AI-powered algorithms are not currently subject to any kind of external and internal regulation. And if and when regulatory frameworks solidify, the compliance requirements may inadvertently lead to rising concentration risk in the industry. The cost of complying with some regulations may simply be too high for smaller developers or users to bear, leaving only a few major large actors in control of large swaths of the AI ecosystem.

Beyond regulatory compliance, failure to protect data can result in meaningful financial expense. [The average data breach in the U.S. can cost companies or other entities more than \\$8 million in remediation expenses](#), with costs rocketing higher for issuers that lack security automation and incident response mechanisms.

Bias: Many AI applications can only be as neutral as the humans that program their decision-making algorithms. A criminal justice algorithm used in Florida, for example, mislabeled African-American defendants as “high risk” at nearly twice the rate of white defendants, as cited in a 2019 Harvard Business Review article. There are other examples of tools such as natural language processors leading to reinforcement of harmful racial and gender stereotypes.

Removing bias from AI implementations requires additional knowledge beyond software and data science, and from career experts who are trained to assess fairness and justice. Morality is not easily evaluated in binary terms. How does the autonomous car facing an unavoidable collision, for example, choose between two crash vectors? Diverse perspectives can help evaluate gray areas, and create equitable processes for managing bias.

We know that these are difficult issues. Governments, companies, and interest groups are all working to find the best path forward. The unfortunate truth is that no one stakeholder has the perfect solution, but as investors, we can learn from the leaders in the space, ask tough questions of our current and prospective portfolio holdings, and encourage progress.



Leveraging Our Investor Toolbox: Engagement

As stated earlier, we are eager investors in a number of companies and bond issuers with leadership roles in the AI space. In our view, companies that can effectively mitigate AI-related risks, and spearhead development of systems and processes that help other companies do the same, stand to gain potential long-term revenue growth and/or avoid meaningful costs over time; these are clearly outcomes we seek for holdings in our equity and fixed income portfolios. Further, corporations and organizations using these technologies can greatly mitigate financial risk through proper and responsible implementation. We engage frequently with management teams and other stakeholders on these issues. We are investors, not AI or computer science experts, so we collaborate extensively with experts in industry and academia to inform our engagement work.

Below, we list some of the ethical AI-related practices and actions we encourage companies and bond issuers to consider during our engagement discussions:

AI Governance:

Add ethics professionals, policies, and review processes to AI product and systems development teams. Too often, developers and users are putting AI products and services into the market before anyone starts asking ethical questions. At that point, it may be too late to fix problems. If ethical questions, asked by ethics professionals, become a part of the development process, we believe that many problems can be avoided before they arise. Such an approach requires a gradual shift away from a “results-at-all-costs” corporate mindset, but that transition can lead to better long-term financial outcomes through risk mitigation and increased brand value.

- We frequently speak with companies about attracting talent with the training and skillsets needed to direct AI systems in responsible directions, and backing up those professionals with sufficient decision-making authority, support teams and other infrastructure so they can succeed. We see Microsoft, for example, as one of the leaders in AI governance practices; it appointed a Chief Officer of Responsible AI, Natasha Crampton, and supports her with ample resources as well as a comprehensive framework and rules for upholding what it refers to as “AI Principles.”

AI Assessment:

Forecast potential risks comprehensively, and disclose those risks through broad and consistent public reporting. Beyond the core benefit of documenting risk as a prerequisite for mitigating that risk, we believe that a steady drumbeat of respected leaders talking about the risks of unethical practices or a lack of diligence will help raise the bar for all AI industry participants. Salesforce, for example, has assessed its underlying AI risks and identified bias as one of the key drivers of “AI error” in its systems and processes. By identifying underlying bias, Salesforce is better positioned to proactively respond to risk before it becomes harm.

Involve heavily impacted communities through representation and other means. Good AI risk assessment can identify vulnerable segments of society that may suffer disproportionate harms from an AI deployment. For example, experts and community leaders have repeatedly highlighted the risks of racially biased results from facial recognition software in recent years. If developers and users of AI incorporate these perspectives, they can avoid the kind of unjust outcomes and backlash that several companies and governments have already experienced.

Be supportive and collaborative with respect to productive privacy laws, not resistant. While privacy laws can create compliance burdens, they can also create regulatory certainty as well as a framework for increasing public trust. Developers and users of “big data” systems should be active participants in collaborating to advance such regulation, recognizing that not all regulation is created equally.

AI Monitoring:

Self-scrutinize policies, actions and results frequently, and report regularly on success and failure to stakeholders. We are particularly focused on asking AI developers and users to provide accurate and transparent reporting to investors; such reports are critical to our ongoing due diligence. This reporting is at its best when it goes beyond standard metrics, and reveals self-reflection on emerging problems along with solutions to address those problems.

Where We Are Today

Artificial intelligence is poised to be one of the key drivers of progress in society for decades to come, and it is already generating tangible positive results in many ways—including in several arenas like energy efficiency that ESG-aware investors find attractive. The AI issues highlighted in this article are complex challenges for governments, investors, companies, and interest groups alike, but finding solutions to them will only benefit all of these parties—investors included—going forward. [B](#)

The views expressed are those of Brown Advisory as of the date referenced and are subject to change at any time based on market or other conditions. These views are not intended to be and should not be relied upon as investment advice and are not intended to be a forecast of future events or a guarantee of future results. Past performance is not a guarantee of future performance and you may not get back the amount invested.

The information provided in this material is not intended to be and should not be considered to be a recommendation or suggestion to engage in or refrain from a particular course of action or to make or hold a particular investment or pursue a particular investment strategy, including whether or not to buy, sell, or hold any of the securities mentioned. It should not be assumed that investments in such securities have been or will be profitable. To the extent specific securities are mentioned, they have been selected by the author on an objective basis to illustrate views expressed in the commentary and do not represent all of the securities purchased, sold or recommended for advisory clients. The information contained herein has been prepared from sources believed reliable but is not guaranteed by us as to its timeliness or accuracy, and is not a complete summary or statement of all available data. This piece is intended solely for our clients and prospective clients, is for informational purposes only, and is not individually tailored for or directed to any particular client or prospective client.