

# Decision-making risks slow down the use of artificial intelligence in business

*The adoption of artificial intelligence promises to bring large benefits to businesses. However, when something goes wrong, all fingers point to the executive who vouched for the technology. The risks are born by decision-makers themselves. **Terence Tse** and **Sardor Karimov** write that this contradiction can explain why AI adoption is moving at a glacial pace in the business world.*

---

Companies have long been interested in exploring and using artificial intelligence (AI) to gain a competitive edge. There is no shortage of media outlets and academics stressing the business benefits that this technology can bring. Such blanket coverage naturally gives us the impression that many companies have already deployed AI internally left, right, and centre, and from top to bottom. Nothing, however, can be further from reality: the speed of adoption of AI solutions by companies has been surprisingly slow. We believe one reason is decision-making risk.

## Decision-making: an important yet understudied risk

Many risks associated with adopting and using AI in business are widely known, as they have been extensively covered in both academic research and consulting reports. These risks range from bias and discriminatory risk to operational and IT risks, from business disruptions to job eliminations. However, a compelling case can be made that a less recognised risk exists, associated with decision-making. During our client interactions, we have observed that executives perceive various forms of risk when purchasing AI solutions. These risks are perceived as very real – and in some instances real enough to discourage them from adopting the AI technology altogether, despite the business advantages that it would bring. To illustrate decision-making risks, we have framed them as questions that the various stakeholders involved in the decision of buying the AI solution would be asking themselves:

### Who among us is taking responsibility if something goes wrong?

It is not uncommon for large companies to benefit from in-house technology assessment teams whose roles in relation to AI involve determining the solution's technical viability and its technological fit with the company's needs. Interestingly, at the point of purchasing, what often matters most is not just how sound the technology is. Instead, it is *who* within the assessment team will be willing to stand up to vouch for the purchase. We have frequently observed that while as a team the members can reach consensus on the evaluation of a solution, as individuals few of them might be inclined to be ultimately responsible for the buy decision. The reason: if down the road something goes wrong and the technology does not work out as intended or expected, fingers would most likely be pointed at the person who "bought" the technology.

The complexity of AI projects makes it difficult to be certain of how the technology will behave once implemented and integrated into the business. Would the AI model work once live data are fed into the system? Would we be able to ensure data security over time? Would the IT infrastructure be capable of supporting the AI solution in the long run? Would the in-house engineering and data science teams be competent to keep the technology going and evolving on a long-term basis? There can be too many questions with too few confident answers.

### Do you know what we are purchasing exactly?

In many companies, business leaders are the ones deciding whether to invest in an AI solution. Many of them are totally dependent on their own data science teams to evaluate and tackle the technological aspects of the solution. In turn, the data scientists have to understand the proposed solution inside-out. They need to know exactly and comprehensively what the company is buying.

Of course, vendors and suppliers can and should help by being completely transparent about their product offering, thereby making it easier for the data science team to understand the technology and its application to the business. Yet, this only partially solves the problem given that many perceived uncertainties remain for the data and leadership teams, including whether the vendor will be around in the long term to provide the necessary support, how scalable the solution purchased is, given the state and nature of the existing IT system, *who* is going to do *what* to maintain the technology going forward, etc. The result: it is often easier and safer for the data science teams – and hence the executive teams – to simply say no to the purchase and avoid all these risks.

### **Can we afford to make the wrong decision?**

Many organisations, especially public ones, have no other choice than to get the technology right the first time. Any purchase decision *must* be turned into a future success. This is because failure to get the technology to deliver what was promised is likely to be construed as managerial incompetence or worse, a misuse of public funds. Under public scrutiny, any outcome short of the anticipated results being *fully* met would potentially lead to reputational if not legal damages.

It must be added that this is not a quandary restricted to public entities. Financial services companies, for instance, can also face the same decision-making risk. Take the example of a bank that is contemplating using AI to cut the costs related to due diligence and meet ever-tightening compliance regulations. It can be imagined that if the AI technology does not work properly, the bank may end up facing costly financial consequences in the form of penalties, compensation, and remedial actions. In view of these, the stake may be too high even though the AI solution could be bringing substantial benefits.

### **Does the technology really work if we cannot figure it out ourselves?**

Behavioural economists refer to “reference dependence” to mean that we assess the outcomes of an action by comparing it with an existing reference point or status quo. This phenomenon can sometimes be observed in the process clients follow to decide whether to buy an AI solution. Many non-tech organisations have employed very talented data and AI scientists, who are tasked with building models to achieve specific business goals.

Yet, in some cases, even with heavy investments of time, money and effort, these in-house scientists were not able to develop an AI model that meets the minimum usability standard or the model only gets an accuracy of 60 per cent when it needs to reach at least 80 per cent to be somewhat useful. Business executives would end up subscribing to the belief that it is simply impossible to get to 80 per cent. Even if a vendor was able to offer a model that demonstrates an 85 per cent accuracy, the scientists would take a dim view and discount such possibility. We should of course beware the vendors’ claims of their technological solutions – there are just too many AI firms making exaggerated claims about the performance of their technologies. Yet, this example illustrates just how powerful this perception can be, which leads to the decision not to buy a technology from an external vendor.

### **Can we justify the spending?**

One of the most commonly used criteria to decide whether to invest in an AI solution is *return on investment* (ROI). Naturally, it is always better to generate as high a return as possible, given a fixed invested amount. So, whereas the cost savings generated from implementing the AI technology in a single department within a company may represent a good return, it would probably be much better to have another department make use of the same AI solution. From the decision-maker’s standpoint, it would be easier to justify the investment with a higher ROI.

In addition, having two sources of return is a safer bet just in case one of them does not deliver the expected financial outcome. The consequence is that executives sometimes only decide to buy an AI solution when two or more departments are willing to commit to its deployment. Investment in AI projects can therefore end up getting delayed, if not cancelled.

Decision-making risks present serious obstacles to the adoption and development of business AI solutions. All too often, our attention on the risks of AI is confined to those risks that relate to such aspects as the AI models themselves, the disruption to business as usual, cybersecurity, etc. By themselves, these risks are often sufficient to discourage companies from adopting AI solutions. However, as discussed here, the risks perceived by decision-makers can actually kill the potential AI projects altogether. This is understandable because, effectively, much of the upside gains of the purchase decision are accrued to the companies, whereas the decision-makers bear the downside themselves. No wonder the adoption of AI can be moving at a glacial pace.



*Notes:*

- *This blog post represents the views of its author(s), not the position of LSE Business Review or the London School of Economics.*
- *Featured [image](#) by [CRYSTALWEED cannabis](#) on [Unsplash](#)*
- *When you leave a comment, you're agreeing to our [Comment Policy](#).*