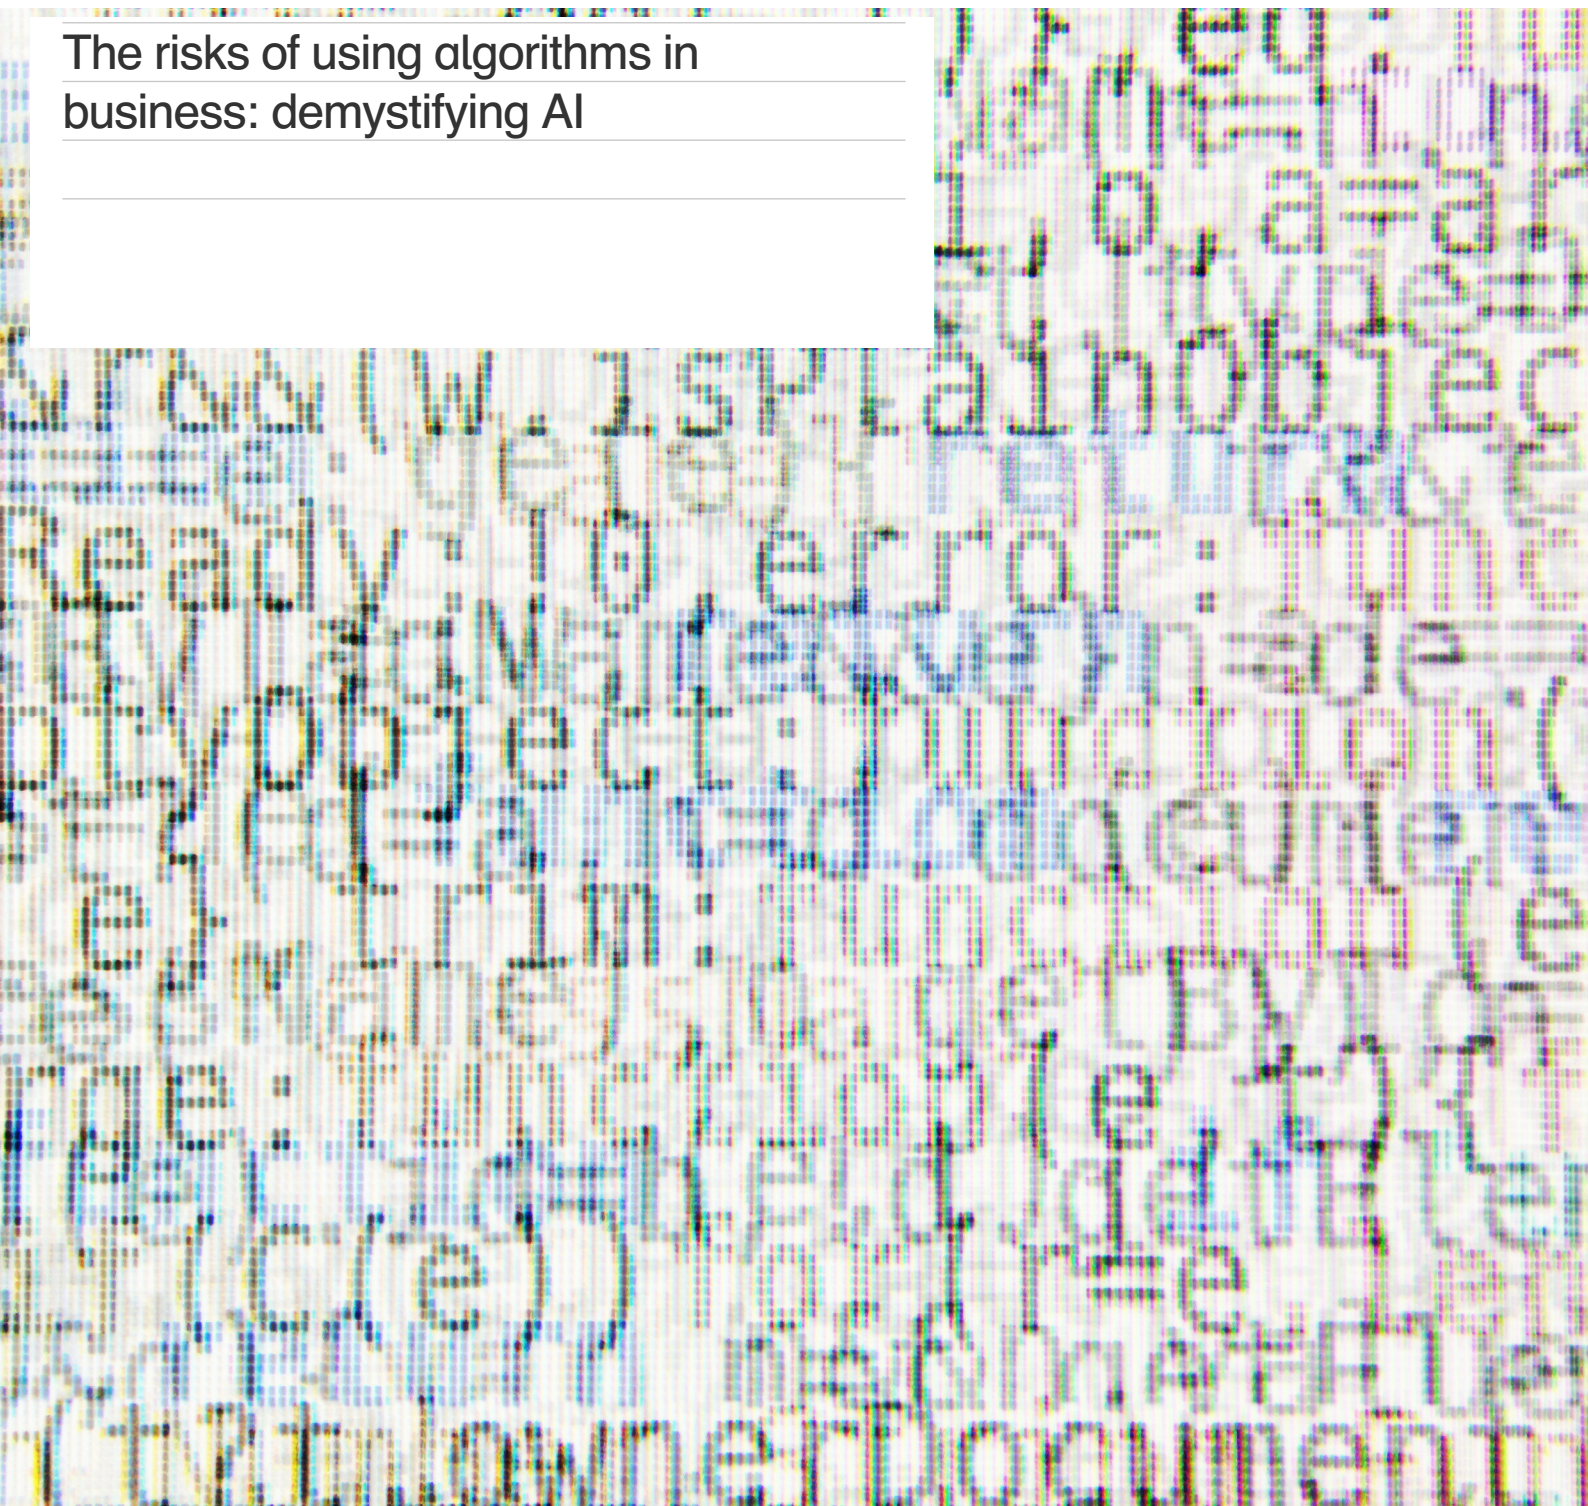


## The risks of using algorithms in business: demystifying AI





**Contact**  
Pascale Déchamps  
Partner

**Today algorithms influence all aspects of our lives, from how much we pay for groceries and what adverts we see, to the decisions taken by health professionals. As these tools become increasingly widespread, they pose new challenges to businesses. In order to begin to demystify algorithms and AI, we ask: what benefits and risks do they bring to the economy?**

Algorithms influence everything from Netflix or Amazon recommendations to online advertising, credit scoring, airline prices and crime prevention. Some commentators even claim that algorithms are a form of ‘artificial intelligence’ (AI), giving the impression that some form of intelligent being is taking decisions that are imposed on us, lending them an aura of neutrality and irrevocability.

The economic and societal consequences of algorithms and AI are profound and take many forms. They affect our purchasing options, which schools our children attend, and the nature of our jobs. From a business perspective, AI comes with regulatory risks; and European and some national regulators are in the process of preparing guidelines for companies, as outlined by the European Commission’s AI White Paper.<sup>1</sup> AI may also pose reputational risks, as is exemplified by recent controversies over the use of algorithms by some companies and governments—for example, the use of algorithms in France and the UK to determine the grades that underpin students’ university applications.

These risks mean that algorithms must be carefully designed, implemented and monitored. Thankfully, economists have a number of tools both to assess the potential harm that could arise from using AI in decision-making processes, and to correct potential flaws.

This article is the first of a series of three *Agenda in focus* articles investigating the risks associated with the use of AI in business. The series aims to answer the following questions.

- How does AI, in its simple and complex forms, work in practice, and how does it affect businesses and society?

- Could AI introduce unfair or even illegal discrimination between individuals or groups, and how can a business ensure that its use of AI does not generate or exacerbate this?
- What are the antitrust risks generated by AI?

In this article, we explore the first question. We explain what AI is, how it is used by firms, and what forms it can take. We also set out the key concepts that are needed to understand more complex economic questions relating to AI.

## What is AI, and are firms really using it?

Many firms use so-called AI or other forms of automated decision-making tools with the aim of improving or accelerating decisions. The applications of AI are numerous and highly varied.

Among other things, AI has been used to fight crime by optimising police patrols, to assist health professionals by improving breast cancer detection from mammograms, and to help sellers with price-setting on online platforms. In our everyday lives, AI is used in voice recognition tools, content filtering on our favourite streaming platforms, and to select the pictures that we see on social media.

Using the term ‘AI’ can be misleading, as it gives the impression that algorithms have a power of prediction or classification that humans do not, thereby potentially leading to a false view that ‘it must be right because the computer says so’. One may get this feeling, for instance, when watching airline ticket prices change from one day to the next for no apparent reason.

AI has been defined by Elaine Rich, author and world-leading academic in computer science, as ‘the study of how to make computers do things at which, at the moment, people do better’.<sup>2</sup> This broad definition carries an important nuance in that it compares machines and humans: AI is about a computer system acquiring human abilities—and perhaps ultimately improving upon them. A notion of temporality can usefully be added to this definition: today’s machines are already much better than humans at doing a wide range of things using vast computing power—AI is about getting machines to do even more. AI also typically involves a degree of interaction between humans and machines.

To make AI a more manageable concept, it can be associated with programs that are not made up of a simple set of step-by-step rules for decision-making, but rather a complicated web of processes that combine

to give a probabilistic result, as if the machine has developed an ‘intelligence’ of its own.

Firms and institutions today use programs and statistical models of varying complexity. Regardless of their sophistication, such tools share common features and risks when decision processes are automated.

## How does AI relate to algorithms, statistical modelling and machine learning?

Computer programs, statistical models and AI are designed to help humans classify information (e.g. automatically classifying emails as spam or, in the classic example used in the field, separating images of cats from images of dogs), match preferences (e.g. in dating websites or song selection algorithms) or predict outcomes (e.g. in weather forecasting or assessing the likelihood that a convicted criminal will reoffend).

## Algorithms and statistical modelling

The most simple forms of such programming involve **simple algorithms** (i.e. human code-based decision-making processes of the type ‘if X is true, then Y, otherwise Z’), as well as **statistical modelling** (i.e. human-designed statistical modelling aimed at explaining a particular phenomenon using factors that influence it, such as classical regression analysis). These two forms of programming can be combined, but in their ‘simple’ form they are *always* designed by humans.

These relatively simple tools can bring productivity gains to organisations, but may also create risks for businesses and society. For instance, this type of program was used in the UK in the 2020 controversy regarding students’ A-level results (see the box overleaf). We will discuss potential risks associated with these models in the second article of series.

## Statistical learning methods

The next level of sophistication involves machines ‘learning’ how to reach the desired outcome (such as classifying information, matching individuals or predicting a result). If the objective is set by a human, the machine decides the best way to achieve it, through **‘machine learning’**.

## The A-levels controversy in the UK

Human-designed programming, such as simple algorithms and statistical models, lay behind the A-levels controversy that caused public uproar in the UK in Summer 2020 (and a similar type of programming was used in France).<sup>\*</sup> It was used to predict what grades A-level students would have earned had their exams not been cancelled due to the COVID-19 pandemic. The objective was to avoid grade inflation that might have resulted from using teachers' predicted grades directly. To that effect, the UK's examination watchdog, Ofqual, assumed that each school's distribution of grades would be similar to that of the previous year. Ofqual used a simple algorithm that took each school's distribution of grades in the previous year and, with only limited adjustments, applied it to the current year's students at the same school. The algorithm used teachers' ranking of students, giving each student the grade of the student of the same ranking from the previous year, but made very little use of the teachers' grading of individual students themselves.

Note: \* Ofqual (2020), 'Awarding GCSE, AS, A level, advanced extension awards and extended project qualifications in summer 2020: interim report', <https://bit.ly/3mXOQvf>. For the French example, see Le Ministère de l'Enseignement supérieur, de la Recherche et de l'Innovation (2020), 'Algorithme national de Parcoursup', <https://bit.ly/3jaZXPL>.

Source: Oxera.

In order to learn, the machine generally relies on a training dataset, a test dataset, one or more techniques, and an objective (such as minimising the prediction error). The process is described in the box below. Machine learning relies on (sometimes complex) algorithms and may make use of statistical modelling. Yet, the form of these sub-programs is not necessarily imposed by humans. Data scientists nevertheless typically understand the precise mechanics of their models and are able to explain the process used by the machine to reach a decision.

However, the machine may also use other methods, such as artificial neural networks or deep learning,<sup>3</sup> which typically become 'black boxes' even for data scientists (in the same way that we may be unable to explain how the human brain reaches a particular decision).

Machine learning is particularly valuable when one is 'data rich and understanding poor'—which means that there is a lot of data that can be used to train the algorithms and little prior understanding about how to take optimal decisions.

Within machine learning, there are different levels of 'machine autonomy':

- **'supervised'**—when the training dataset contains data labels identified by humans (e.g. when pictures of cats are identified and labelled as such by humans in the training dataset);
- **'unsupervised'**—when the training data is not labelled and the machine needs to find patterns by itself. For instance, the machine may use a survey to identify groups of users that share common characteristics;<sup>4</sup>
- **'reinforcement learning'**—when the machine learns through independent trial-and-error exploration, and amends its future decision-making process to improve optimisation.

Machine learning techniques are increasingly used by petrol stations for pricing. Indeed, there are AI companies such as the Danish a2i systems and UK-based Kalibrate that offer AI pricing software to optimise dynamic petrol pricing. Initial training is based on historical data (such as past transactions, competitor prices and other market conditions), after which prices are set by taking into account 'real-time' information (such as competitors' current

prices, the weather, and traffic conditions). The resulting transactions are then in turn fed back into the system and used to re-optimize the algorithm. These programs therefore make use of supervised learning in combination with a reinforcement mechanism.

The use of such petrol pricing algorithms can benefit both petrol stations and consumers, as it reduces the cost of managing prices and increases market efficiency. Petrol stations pass on lower costs and faster price adjustments to consumers who pay lower prices. However, competition authorities have voiced their concern that such pricing algorithms may be used as an instrument of collusion to sustain supracompetitive petrol prices.<sup>5</sup> Future articles will discuss how problems such as this may emerge.

## What does all of this mean for business and society?

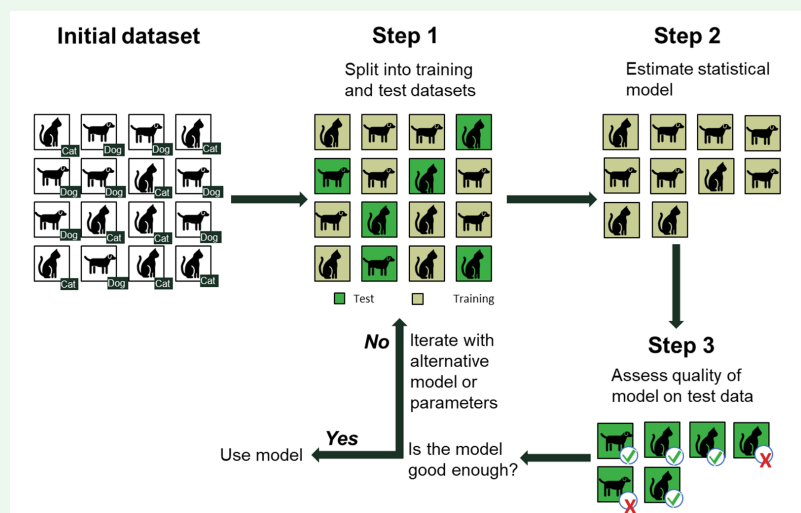
Computers have been assisting human decision-making for a long time. Most of the underlying techniques have been thoroughly used by humans in a wide range of applications. Forecasting the weather

## Illustration of the concept of machine learning

A machine learning algorithm is fed a dataset (such as pictures of cats and dogs), at least one technique, and one objective. The algorithm estimates a model using the training dataset. The complexity of such models can vary widely, from simple linear equations to complex neural networks. If the data is labelled, the machine automatically 'knows' whether a picture represents a cat or a dog, and therefore follows a supervised learning process. If the data is *not* labelled, the machine instead follows an unsupervised learning process.

In the test dataset, the quality of the initial model is assessed in light of the objective that was set (e.g. predicting whether a picture accurately represents a cat or a dog in a supervised learning process). If the objective is reached in the test dataset, the model can be used. If the results are not satisfactory, a new mathematical model is estimated on the training dataset until the objective specified in the algorithm is reached in the test dataset.

## Illustration of machine learning in a supervised learning problem



Source: Oxera.

or changes in house prices has been undertaken with increasing sophistication for decades.

The recent change has been the exponential rise in computer power combined with the collection of enormous amounts of detailed data on all parts of our economies and lives. For instance, it is estimated that there will be 44 zettabytes (44 x 10<sup>21</sup> or 44,000,000,000,000,000,000,000 bytes!) of data by the end of 2020, of which 90% was created since the start of 2019.<sup>6</sup>

This change allows programs to analyse more complex data using more complex techniques at a lower cost. In order to streamline processes, companies and governments alike have invested in designing programs to help their decision-making or customer interaction at a scale that would have been unimaginable a few decades ago. The way in which algorithms play roles in our lives would have been unthinkable in the recent past.

The widespread use of algorithms has changed our lives and the way businesses operate. They affect which candidates are hired, how prices are set, and which adverts are shown to potential consumers. The rise of AI and machine learning offers unprecedented opportunities for businesses and other organisations.

However, the use of AI does not come without risks and challenges. As mentioned above, a key risk is firms' ability to circumvent competitive market pricing rules in a way that is not possible for humans. Another relates to algorithms' role in worsening or preventing discrimination in society.

There is also a wider concern regarding AI's ability to undermine humans' core values by not taking properly into account the ethical and societal impact that some of these programs may have, beyond their immediate application.

Governments and regulators are increasingly paying attention to these risks. Some regulators already have a framework to account for algorithms in business,<sup>7</sup> while others are evolving. For instance, at the EU level, the Digital Service Act (DSA) is likely to introduce some new tools for the regulation of AI and algorithms. In France, ARCEP, the telecoms regulator, considers that algorithms' 'explainability'<sup>8</sup> and accessibility to regulators is an essential part of platform regulation.<sup>9</sup>

This means that companies are facing new risks to their operations. In the forthcoming articles in this series we will investigate two prominent challenges that have gained public attention:

- the risk of exacerbation of illegal discrimination by algorithms (but also why algorithms offer a unique opportunity to reduce discrimination);
- the antitrust risks created by pricing algorithms.

By studying these matters, we will shed light on some of the main challenges that businesses face in their use of AI, but also how economists can help to overcome these challenges.

## Contact

[ambroise.descamps@oxera.com](mailto:ambroise.descamps@oxera.com)

Ambroise Descamps

[pascale.dechamps@oxera.com](mailto:pascale.dechamps@oxera.com)

Pascale Déchamps

[gareth.shier@oxera.com](mailto:gareth.shier@oxera.com)

Gareth Shier

[timo.klein@oxera.com](mailto:timo.klein@oxera.com)

Timo Klein

<sup>1</sup> European Commission (2020), 'White Paper on Artificial Intelligence – A European approach to excellence and trust', 19 February, <https://bit.ly/3cBUkY5>.

<sup>2</sup> Rich, E. and Knight, K. (1991), *Artificial Intelligence*, McGraw-Hill Education.

<sup>3</sup> An artificial neural network is a highly flexible (but data-intensive) type of machine learning that is inspired by the functioning of the human brain. It works by creating multiple layers of input and output nodes (loosely comparable to the interconnected neuron nodes in a human brain) and optimising the links between these nodes based on the training data that it 'sees'. Put simply, we talk about 'deep learning' when we use artificial neural networks that consist of several layers.

<sup>4</sup> For an application, see Oxera (2018), 'Android in Europe, Benefits to consumers and businesses', prepared for Google, October, <https://bit.ly/36dmuHv>.

<sup>5</sup> *Wall Street Journal* (2017), 'Why Do Gas Station Prices Constantly Change? Blame the Algorithm', 8 May, <https://on.wsj.com/3mYhPiy>.

<sup>6</sup> BBC (2017), 'Big data unites the study of stars with cancer research', 13 March, <https://bbc.in/3n0C675>.

<sup>7</sup> European Commission (2020), 'White Paper on Artificial Intelligence – A European approach to excellence and trust', 19 February, <https://bit.ly/3kSz9np>.

<sup>8</sup> 'Explainability' is increasingly used to mean that an algorithm's designers should be able to describe (in a non-technical manner) the way in which an outcome has been produced given the models and dataset used. For instance, see European Commission (2020), 'White Paper on Artificial Intelligence – A European approach to excellence and trust', 19 February, <https://bit.ly/3cBpkHF>, p. 5.

<sup>9</sup> Arcep (2020), 'Arcep's contribution to the public consultation on the DSA Package and the New Competition Tool'.