Artificial Intelligence, Big Data and the Human Mind: A Study on The Effects of New Technologies on Students' Decision-Making Process

Jacopo Vescovi

September 6, 2018

Master Thesis

Copenhagen Business School

Master in Management of Innovation and Business Development

Supervisor: Keld Laursen

Student Number: 113813

Abstract

Recent developments in the fields of machine learning and data management are drastically changing the way companies and individuals consume information necessary for decision-making. However, the effects that exposure to these new techniques has on the human decision-making process is still unclear, in spite of the benefits that wider knowledge on the topic could bring to businesses.

The aim of this thesis is to advance the current state of research on the matter, by investigating the effects that taking courses on Big Data or Artificial Intelligence has on students enrolled in master's degrees at Bocconi University or Copenhagen Business School. A review of the current literature led to the identification of three tendencies part of the students' decision-making process worth considering: their reliance on data rather than opinions, their preference for internal rather than external innovation, that is to say Not-Invented-Here Syndrome, and their propensity for delegating a task to a machine rather than doing it in person.

The students were surveyed before and after half of them took part in courses concerning Big Data or Artificial Intelligence, testing their attitudes with vignette scenarios. Their responses were then analysed with a differences-in-difference approach.

The results show that exposure to Big Data and Artificial Intelligence increases data literacy and causes the development of a unique framework for approaching different decision-making situations. On the other hand, there seem to be no significant effects on Not-Invented-Here Syndrome and on the tendency to delegate to machines, as these phenomena are not altered by the attendance of courses on Big Data and Artificial Intelligence.

The implications of these outcomes are then discussed, showing how they can provide useful information and outlining the positive effects of the courses considered for students and companies alike. In particular, the courses are shown to increase students' ability to use data without creating a bias for automation or a blind trust in machines and software, as respondents remain aware of the importance of human work.

The final part of the thesis presents the findings' implications for business and academia, along with interesting related topics that future research may consider worth investigating.

Contents

1 Introduction

This work will deal with the effects that the exposition to and use of new technologies, such as Big Data and advanced analytics powered by AI, have on the human mind.
These effects are probably too various and complex to be investigated in their entirety in the present research, as they often involve changes that concern not only the individual, but also organisations and society (Brynjolfsson & McAfee, 2014).
As a matter of fact, while Artificial Intelligence and Big Data have been the subject of much research in recent times, the current state of research tends to focus on the technical aspects and the effects that these new technologies can have on society along with the way data analysis is conducted.
On the other hand, few articles concerning the consequences that AI and Big Data may be having on the human mind can be found, and most of the published works are the results of theoretical, rather than experimental, analysis.
The issue is especially relevant in the field of management, where it could be caused by the novelty of these technologies, so that it is merely too soon for a comprehensive body of research to have developed, or by the complexity of the topic, that makes them hard to identify and experiment with.
This thesis aims to reduce this gap, by using experimental data to explore the validity of hypotheses generated on the basis of the current body of literature.

The aspects of the human mind that could be analysed are vast. Here, the focus will be on the decision-making process, more specifically on three tendencies that decision makers may exhibit in their thought process while making a choice and how likely they are to be changed.
The reduced scope of this research derives from the need to increase the feasibility of an experimental study that could be performed on a sample of students, who could have been overwhelmed by a research testing dozens of different personal features and tendencies.
Additional considerations affected the final judgement on which topics were the most promising to explore, here described more in detail.
First of all, a better understanding of how the decision-making process is affected by the technologies that are being developed and implemented in thousands of companies could help businesses make better choices and harness the benefits these technologies can offer.
Secondly, knowing whether there are any potential pitfalls in the way decisions are taken with these technologies could help companies to prevent them from doing any harm. Sub-

sequently, not only more informed decisions could be taken by decision-makers, but also nefarious choices could be more easily avoided.
Thirdly, this thesis can contribute to the present state of the research by providing an experimental corroboration or confutation of different theories, whence a deeper understanding of the relationship between new technologies and decision-making processes could stem.
Lastly, this research can be helpful for universities offering courses concerning Big Data and AI, that expose students to these kind of technologies. By being aware of the effects these courses can have, institutions can improve their curricula and verify the validity of their learning objectives, along with the ability of the courses to achieve them.

The three potentially affected tendencies examined in this thesis emerge mostly from the work of Brynjolfsson and McAfee (McAfee & Brynjolfsson, 2012), and in minor part from the work of Migliore (Migliore, 2017), who describe the features required by companies and individuals to fully harness the latest analytics technologies.
Jointly testing these features and how they change after exposure to Big Data and AI is akin to testing an overarching research question that can be defined as follows:

Does exposure to advanced analytical tools and topics relying on Big Data and Artificial Intelligence modify human decision-making tendencies other than increasing data literacy?

This can be further declined in 3 sub-questions, that reflect the tendencies tested:

Sub-Question 1 *Does exposure to advanced analytical tools and topics relying on Big Data and Artificial Intelligence increase a person's tendency to rely on data rather than opinions, or their data literacy?*

The fact that courses on how to use data increase data literacy has been repeatedly observed in prior research. Therefore, the main objective of this sub-hypothesis is the confirmation of previous works as a part of the overall study. Nonetheless, the treatment in this study includes more advanced courses than the ones considered in previous research (Kippers, Poortman, Schildkamp, & Visscher, 2018; Dunlap & Piro, 2016), and besides focusing on the latest data trends it spans an entire semester instead of a couple days. Therefore, the contribution that the present thesis can give to the debate on this topic is not negligible.

Sub-Question 2 *Does exposure to advanced analytical tools and topics relying on Big Data and Artificial Intelligence increase a person's tendency to incorporate external innovation rather than developing technology in-house?*

Sub-Question 3 *Does exposure to advanced analytical tools and topics relying on Big Data and Artificial Intelligence increase a person's tendency to delegate tasks to a machine rather than doing them in person?*

Starting from the analysis of the work by Brynjolfsson, McAfee and Migliore (McAfee & Brynjolfsson, 2012; Migliore, 2017), the aforementioned hypotheses were formulated. Each of them was tested through two surveys with a scenario, for a total of three scenarios per survey. The responses were then analysed through the analysis of the differences in the answers provided by the survey respondents in the two questionnaires. A brief overview of the three features embedded in the scenarios, as first described by the authors, follows.

The first feature is the ability to rely on data rather than opinions, especially if the latter are used systematically, a phenomenon known as HiPPO, or Highest Paid Person Opinion. This work tests the tendency to rely on data rather than previously held opinions, that may be affected by external actors.

The second characteristic is the capacity to integrate outside innovations within the company, thus avoiding the development of a bias for in-house innovation, also known as Not-Invented-Here Syndrome. This work tests the tendency to rely on external, rather than internal, innovation.

The third tendency is not directly described by the authors, but is an original development of a previous work by Hitt and Brynjolfsson (Hitt & Brynjolfsson, 1997) about the effect of IT on the degree of delegation within companies. This work tests the tendency to delegate complex tasks to a software rather than doing it in person.

To explore how the exposition to AI and Big Data technologies affects the aforementioned tendencies, the aptitudes of a sample of business students were tested before and after some of them attended courses on these technologies. The changes in the tendencies of the students were then examined, comparing the differences between the group of students who attended AI or Big Data courses against the group of students who did not, similarly to what is done in difference-in-differences studies.

This method permitted to determine whether the behavioural changes exhibited by the stu-

dents who attended AI and Big Data courses were the results of their enrollment in these courses or were merely due to chance.
As exposure to data through teaching has already been proved to increase data literacy (Dunlap & Piro, 2016; Kippers et al., 2018), that is to say the ability to rely effectively on data when available, the primary aim of this research is not that of proving the same phenomenon again. The objective of this work is instead to test whether exposure to new technologies concerning algorithms and data does affect people's mindset in other ways other than increasing their ability to correctly analyse data.

This thesis is thus organized as follows:
In Chapter 1, an introduction to the research is provided.
In Chapter 2, a thorough review of the literature concerning the topics needed for a proper understanding of the present work is given, and the theoretical basis of this research, along with the current state of research on the phenomena that are the subjects of this analysis, are presented.
In Chapter 3, the methods adopted for this research are illustrated, including the research design and the techniques used to gather experimental data.
In Chapter 4, the methods used for the analysis of the collected data are briefly exposed, along with a model used for the analysis of the data.
In Chapter 5, the results of the analysis are presented and commented on their more immediate and technical aspects.
In Chapter 6, the results are discussed, taking into consideration the relevant literature and exposing the likely causes and implications of the findings.
In Chapter 7, conclusions are drawn, and implications for business and research, along with the limitations of this study, are presented.
Appendix A presents additional information obtained from data analysis which is only briefly referred to in this thesis due to its lower relevance to the discussion.

2 Literature Review

The framework of this thesis is based on 3 main concepts: Artificial Intelligence, Big Data, and Decision-Making Process.
While the third concept is easily given a definition and a clear description (Provost & Fawcett, 2013), the first two elements are harder to define, either because their form is new and constantly changing, or because they are labels attached to a heterogeneous set of new technologies (Provost & Fawcett, 2013).
Moreover, Artificial Intelligence and Big Data are often co-dependent, as recent machine algorithms need vast amounts of data to be trained (Kelly, 2017), and Big Data are celebrated because they allow advanced analysis previously not feasible (Boncea, Ionut, Smada, & Zamfiroiu, 2017; Chen & Storey, 2012).
Hence, the need to provide a clear outlook on the main features of these ideas and the way their development has taken place.
In the following sections, these 3 concepts will be examined more in depth and presented through an analysis of the relevant literature.
In regard to the Decision-Making Process, this review will be focused on Data-Driven Decision-Making, as it is more closely related to Big Data and Artificial Intelligence, which are means to provide data and information whence decisions can be taken.

2.1 Artificial Intelligence

The exact characteristics of Artificial Intelligence have been hard to pin down consistently, as due to the fast progress in the field what was considered the most appropriate expression of AI twenty years ago is different to what it is today.
In the 1990s, the most famous expressions of AI were chess-playing computers like DeepBlue (Barrat, 2015). Today machine learning, a branch of statistics, is the field that most represents AI (Dickson, 2017). Its applications power cars that can drive themselves, provide advanced speech recognition, and are able to autonomously create music and art, although under human supervision (Galeon, 2017; Pogue, 2018). Moreover, Artificial Intelligence tools are used for more mundane but more impactful tasks like fraud detection and medical diagnosis (Baesens, Bapna, Marsden, Vanthienen, & Zhao, 2016) and are even starting to replace accountants (Issa, Ting, & Vasarhelyi, 2016).
On the other hand, many of the tools that were once regarded as the bleeding edge of Arti-

ficial Intelligence have now become commonplace and are not even considered AI anymore, because they have lost their novelty and the ability to elicit a feeling of wonder, like spam filters. As a matter of fact, when we think of AI, the first examples that come to mind are futuristic robot designs and the promises of machines that can act and think on par with humans (Kurzweil, 2006).

Despite the apparent science-fiction character of this vision, many technologies are already capable of performing certain tasks in specific fields better than humans can. In many strategic games, computers can beat top humans, as DeepBlue (Roberts, 2016) and AlphaGo (Silver et al., 2017) have showed in chess and go respectively.

Despite this progress, in other fields more similar to the popular conception of Artificial Intelligence, computers are still lagging far behind the best humans, although their performances have recently greatly improved.

Speech recognition and translation, in particular, are now routinely and reliably used in the form of personal assistants like Apple's Siri or services like Google Translate (Boudreau, 2017). What was once the wondrous state of the art is now part of everyday life, and AI is already, in one form or another, in many of the tools that are used for work and leisure alike, being entrusted with high-level tasks like texting on one's behalf.

The enormous variety of tasks that machines are able to accomplish derives from the original idea of "universal machine" that Alan Turing envisioned in the time between the two world wars (Turing, 1936), a machine capable of executing any conceivable set of instructions.

Turing was also one of the pioneers of the concept of "thinking machines" that could execute tasks as well as human beings, becoming indistinguishable from them and thus reaching Artificial General Intelligence, the definition of machine intelligence on par to that of humans. He most notably conceived the Imitation Game, later developed into the Turing Test, to evaluate the similarity between machine and human behaviour (Turing, 1950). Although the test has been heavily criticized for not correctly evaluating or representing the advancements in AI technology (Norvig & Russell, 2009), it provided one of the first tools to assess the capabilities of computers.

Since then, AI research has gone through alternate phases of hype and optimism, followed by pessimism stemmed from broken promises and funding cuts. The periods of time that saw research stagnate are known as "AI Winters", and they usually followed a downturn in the perception of the technology by government and venture capitalists, that as a consequence cut funding for the field and starved basic research (McCorduck, 2004).

In the present day, the world is experiencing a new AI spring, driven mainly by machine learning and in particular by neural networks, which allow computers to accomplish tasks that were previously thought to be the hardest for a machine to execute, like driving a car

(Waters, 2017). Venture capitalists and governments are racing to fund machine learning research and applications, and the US and China are engaging in a AI arm race (Lucas & Waters, 2018; Schechner, MacMillan, & Lin, 2018).
Moreover, the ability of neural networks to perform better than humans at image-recognition and prediction tasks has driven the adoption of AI in industries of critical importance and value like energy and transportation infrastructure, finance, trading and commerce (Barrat, 2015).
The imbuing of so many activities with Artificial Intelligence has brought to a revolution in how companies and employees work with and relate to computers (Brynjolfsson & McAfee, 2013; McAfee & Brynjolfsson, 2016; Shook & Knickrehm, 2017; The Economist, 2017a). Software and algorithms are now complementary or even essential in many decisions that were once the domain of human intuition. This is due to the specific advantages that machines have over humans when performing certain tasks, and employing Artificial Intelligence where these advantages provide an edge can boost a company's value and returns (The Economist, 2018).
As a matter of fact, two of the main advantages that machines have over humans is the sheer quantity of data they can handle and the rigorousness of their decision process. The first element has enabled the staggering success and popularity of the Big Data movement (Boncea et al., 2017; McAfee & Brynjolfsson, 2012), while the second has reaffirmed the need for data-driven decision-making processes in business, as it can improve performance and boost value in a company (Brynjolfsson, Hitt, & Kim, 2011; Brynjolfsson & McElheran, 2015).
However, these new analytics technologies create not only benefits, but also challenges. As the accuracy of algorithms grows, the risk for companies and individuals to incur into automation bias increases (Parasuraman & Riley, 1997). Automation bias can be a serious danger to the decision-making process of a company, as over-relying on machine aids for decisions can lead individuals to underperforming (Skitka, Mosier, & Burdick, 1999).
Beyond decision-making, which happens even at very low hierarchical levels in companies that have invested heavily in IT (Hitt & Brynjolfsson, 1997), AI also has the potential of revolutionizing the way whole industries are organised. By learning to execute tasks that at the moment require human employees (Brynjolfsson & McAfee, 2014; The Economist, 2017b), it will force them to adapt to a sharply altered job landscape. The steady advances in AI technology will likely require people to change their tasks and update their skills regularly over their lifetime (Shook & Knickrehm, 2017).
This change will not necessarily impoverish the current jobs performed by humans, as the tasks that are being automated are frequently the most repetitive and boring, and AI advancements could have the ultimate effect of letting human employees focus on more creative

activities, creating new kinds of jobs (Dickson, 2017; The Economist, 2017a).
Thus, AI could not only change the present way work is done and decisions are taken, but also define what kind of skills and behaviour will be needed for the companies and jobs of the future.

2.2 Big Data

Unlike AI, the Big Data phenomenon is quite recent. It has become the latest fashion among tech-savvy companies and CEOs, and recommendations on how to use Big Data to empower business have abounded on both magazines (McAfee & Brynjolfsson, 2012; Wessel, 2016) and peer-reviewed publications (Baesens et al., 2016; Gopalkrishnan, Steier, Lewis, & Guszcza, 2012; Migliore, 2017).
The suggestions of these authors are as varied as the interpretations of the term "Big Data" itself, which originally started to be used as a reference to datasets so large that needed supercomputers to be processed (Manovich, 2001). The power and calculation speed of supercomputers, however, never cease to grow, as does the amount of data they can process. The original definition of Big Data is thus quite vague.
Nowadays, Big Data differ from traditional large datasets in respect to different aspects, summarised by the 3Vs of Big Data: Volume, Velocity and Variety (McAfee & Brynjolfsson, 2012; Gopalkrishnan et al., 2012). Lately, more features have been associated with Big Data to distinguish it further from previous types of data, like Veracity and Value (Fan & Bifet, 2013; Demchenko, Grosso, & Membrey, 2013). These are however unnecessary to present a comprehensive picture of the implications of Big Data in regard to business intelligence (Kimble & Milolidakis, 2015).

The original 3Vs are here exposed in more detail:

Volume The most cited characteristic of Big Data is probably the immense volume of data that are now available to companies, and that are order of magnitudes bigger than even a decade ago (Kimble & Milolidakis, 2015).
This kind of volume permits a granular and highly detailed measurement of people and processes, that results in new services and products, along with predictions that rival those of human professionals (McAfee & Brynjolfsson, 2012).
However, at the same time the efficacy of traditional statistical techniques of analysis can be

undermined by this new phenomenon, as they are not appropriate for the handling of such massive databases (Beer, 2016). Moreover, such an abundance of data can often hide the fact that important information is missing (Gopalkrishnan et al., 2012).

Velocity Velocity pertains to the fact that it is possible to obtain streams of data in real-time and with high precision (Kimble & Milolidakis, 2015), from mediums such as smart-phones, which provide accurate geolocation.
Companies that are able to harness this feature can become more agile than their competitors, and achieve a competitive advantage (McAfee & Brynjolfsson, 2012). This means that there is a necessity for tools that can process rapidly changing information, and provide meaningful insights to decision-makers (Boncea et al., 2017; Demchenko et al., 2013).

Variety Variety concerns the multiple and different sources from which data can originate, and the different formats it is represented in (Gopalkrishnan et al., 2012). As more and more activities are digitized (Beer, 2016) the complexity of the information that is gathered and of the databases that must be built increases. In particular, Big Data can be collected in structured, unstructured, semi-structured, and mixed formats, which need new forms of storage (Demchenko et al., 2013). Therefore, before these data can be utilized by companies to create value, they need to be given a meaning by further processing and analysis (Kimble & Milolidakis, 2015).

The aforementioned characteristics are inducing a change in the way knowledge is acquired, processed and organised (Boyd & Crawford, 2012).
This is a reason of concern for many. As some authors have pointed out (Mayer-Schonberger & Cukier, 2014), vast amounts of data may make it possible, for the purpose of forecasting, to ignore causation. Instead of focusing on unearthing the cause of a particular user trend or feat, Big Data analytics can explore the correlations between known user traits and unknown user preferences to determine the latter, usually with a high degree of overall precision.
Such precision is higher at the aggregate rather than at the individual level. In fact, while it may be hard to forecast what product a particular individual desires, by creating user profiles for millions of people, like Google does, on average the predictions are "good enough". Indeed, good enough predictions and data can be the most appropriate when decisions must be taken in real time or if there are high volumes of users that can thus provide high returns even if the data fall short of adequately describing some of them (Davenport, Barth, & Bean, 2012).

While this technique has been showed to provide useful insights (Mayer-Schonberger & Cukier, 2014; McAfee & Brynjolfsson, 2012; Migliore, 2017), relying heavily on correlation and predictive models rather than on the analysis of causation for business decisions and practices can be dangerous (Ekbia et al., 2015). While data have the appeal of objectivity, and allegedly "speak for themselves" (Mayer-Schonberger & Cukier, 2014), many authors have pointed out that the interpretation of data is an all-too-human process, subject to different flaws and bias (Boyd & Crawford, 2012; Kimble & Milolidakis, 2015). The tendency of humans to see patterns, even where there is none, can be problematic in the face of vast repositories of data, which offer connections radiating in all directions, not always reliable (Boyd & Crawford, 2012). In this context, spurious correlations can be seen and interpreted as real, as distinguishing the spurious from the authentic among thousands of potential correlations becomes harder.

As a consequence, business decisions with far-reaching consequences may end up being taken on the basis of feeble correlations, absent proof of any causation between two variables. Decision-makers must be aware of the risks that Big Data brings to their decision-making process, especially since data without context may provide accurate, but meaningless, answers (Kimble & Milolidakis, 2015).

Moreover, data itself could be biased, because of a bias in the gathering of the data or in its preparation for analysis (Newell & Marabelli, 2015). Especially now that extensive data are available on many topics, there is a tendency to consider a sample as representative of the population based only on the fact that there are millions of observations. This is not always true, as where the data are gathered from is important for considering the set of observations and their relationship to the general population (Boyd & Crawford, 2012), and companies could compromise their products by taking decisions regarding a certain population based on the wrong data.

This trend highlights another problem: data is often analysed "as is", without critically examining individual observations or the gathering process, mainly because of the vastness of the datasets that makes it practically inconvenient. However, this often results in underrating the importance of having good data quality (Baesens et al., 2016), and increases the risk of taking wrong decisions by blindingly trusting data.

Despite these challenges, if Big Data and Big Data thinking are correctly integrated in the decision process of a company, the benefits can be considerable.

Companies can benefit first of all from increased innovativeness in the products and services they offer thanks to new insights previously unobtainable (Baesens et al., 2016; McAfee & Brynjolfsson, 2012).

Secondly, Big Data can greatly improve knowledge management and organizational agility.

A firm can become more aware of its environment and the opportunities and threats that it presents, using this awareness to achieve a strategic advantage. As a confirmation of this advantage, Big Data has been shown to improve business value by facilitating the acquisition of supply chain and marketing knowledge in European firms (Côrte-Real, Oliveira, & Ruivo, 2017).
To successfully harness these benefits, companies need to change and continuously realign work practices, organizational models and stakeholder interests (Günther, Rezazade Mehrizi, Huysman, & Feldberg, 2017). This is probably the biggest impact that Big Data can have on a business, as it requires the organization to change and adapt, evolving its decision-making practices towards a more data-oriented, but critical, approach.

2.3 Data-Driven Decision Making

Data-driven decision making, sometimes abbreviated as DDDM, refers to the process of basing decisions on data analysis rather than exclusively on intuition and gut feeling (Provost & Fawcett, 2013).
This notion and its applications are closely related to the use of AI, data analytics, and Big Data that business makes, as these technologies are utilized to gather, categorise, process and visualize information based on external data, in order to take a decision (Provost & Fawcett, 2013).
Data-Driven Decision Making has been applied to different fields, more so lately possibly because of the ever-increasing availability of data over a broad range of topics, that have allowed more companies to exploit it, often by increasing their IT investments and learning opportunities (Brynjolfsson & McElheran, 2016). The concurrent development of data science has also contributed to shifting more companies towards basing their decisions on data (Provost & Fawcett, 2013).
In this regard, the Big Data movement is a strong promoter of data-driven decision making, as many view data as the ultimate source of objectivity, able to "speak for themselves" (Mayer-Schonberger & Cukier, 2014).
As advancements in computing have contributed to the development of easier ways to analyse data, decision-making based on data has become commonplace (Brynjolfsson et al., 2011), and has thus been the subject of research of multiple studies (Brynjolfsson et al., 2011; Brynjolfsson & McElheran, 2015; Mandinach, Honey, & Light, 2006).
It is generally accepted that data driven decision making is not an all-or-nothing approach (McAfee & Brynjolfsson, 2012), but resembles a spectrum on which companies and individ-

uals can position themselves, deciding to which degree they want to trust data and analytics rather than human experience and insight (Provost & Fawcett, 2013). And while data-driven-decision making adoption has tripled in only 5 years in the US, certain types of companies are more likely than others to use it to take decisions (McElheran & Brynjolfsson, 2016).

As a matter of fact, the adoption of data-driven decision making by a company is affected by various factors, such as the investment in and sophistication level of its IT, its openness to new managerial practices and the educational level of its employees (McElheran & Brynjolfsson, 2016).

These two last characteristics, in particular, make master students an interesting sample to consider, as studying new managerial practices during higher education could make them more likely to experience DDDM. Thus, we can comfortably assume that they know what DDDM is and can adopt a critical and rational approach when exposed to it, making it easier to test their attitudes and reactions.

The unevenness in the adoption of data-driven decision-making highlights the fact that it may not be the best solution for all problems or for all industries, as different issues can arise. First of all, managers could be overseeing environments that do not permit reliable data collection, or where data are scarce (McElheran & Brynjolfsson, 2016). Furthermore, a company may not have the right organizational requirements and skills to harness the benefits of data-driven decision-making.

Moreover, managers themselves are not immune from biases in the interpretation of data, which can make them reach wrong conclusions and act on that basis, resulting in dangerous consequences for the company. This danger is aggravated by the fact that curated graphical data displays do not reduce the risk of bias (Hutchinson, Alba, & Eisenstein, 2010), meaning that simply improving data exposition is not enough to offset this issue.

Despite these challenges, however, the firms that have adopted data-driven decision making have seen considerable benefits from it. Adoption of DDDM by a company has been showed to be linked to improved company performance, causing an increase in productivity of up to 6% in publicly traded companies (Brynjolfsson et al., 2011). A similar increase in productivity has also been found in manufacturers (Brynjolfsson & McElheran, 2015), as long as they are able to implement these new practices, a skill which they seem to be getting better at (Brynjolfsson & McElheran, 2016).

Moreover, data-driven decision-making adoption is a top priority among finance leaders within companies, as it helps them drive growth and earnings (The Hackett Group, 2013).

Looking at these benefits, the importance for managers to be exposed to DDDM practices becomes clear, as learning about DDDM is one of the major drivers behind its adoption (Brynjolfsson & McElheran, 2015). As a consequence, the treatment sample in this study, after being exposed to similar practices, should be expected to present a behaviour skewed towards relying on data to take decisions more than before.

2.4 AI, Big Data and Decision-Making: Working with Machines

All of these elements have been increasingly studied and debated in recent years, although the backgrounds for these concepts have been laid out as early as in the 30s (Turing, 1936). Nowadays, articles about the effects of Big Data and Artificial Intelligence on companies and individuals appear both in peer-reviewed journals (Brynjolfsson & McElheran, 2015) and in other types of media such as magazines (The Economist, 2017a) and newspapers (Lohr, 2012).

These effects are present because AI and advanced tools for analytics, despite major progresses in recent years (Kiulian, 2017; Reed et al., 2017), cannot work autonomously, and still need human labour to provide insightful results.

As a matter of fact, the tasks that machine systems can execute have no sense or meaning if one considers the human and the machine as units in isolation. These are instead agents who execute subtasks obtained by decomposing an overall task and distributing the various pieces among humans and machines (Hoc, 2000).

The two systems are co-dependent from each other, and if humans influence machines because they design them, similarly machines can influence humans by proposing and rewarding a certain type of behaviour that can accomplish the overall task in the most efficient and effective way possible. This trend is embodied in the continuous development of capacities and skills that new tools require from the people who wield them, and the transformations that businesses have to go through to adequately harness innovation (Brynjolfsson & Hitt, 1998).

In the case of computers, very few other tools have required such extensive adjustments in the behaviour of workers and companies. In business, computers have allowed humans to analyse and process large amounts of information, a task which humans are not particularly good at (Issa et al., 2016) but which is important for predictions and for the management of great quantities of products that a global economy makes possible. This has resulted in the centrality of the computer in the business environment, the invention of cubicles, and employees' mansions often reduced to data entering (Habegger, 2005). However, computers have also had less negative effects, like increasing the degree of decentralized authority in companies and the reliance on workers' skills and human capital (Hitt & Brynjolfsson, 1997). Moreover, despite the problematic aspects of adopting computers in business, companies

invested in IT because they could see that it brought higher financial results (Dos Santos, Peffers, & Mauer, 1993; Mahmood & Mann, 2015; Weill, 1992).
However, there is still a lack of convincing evidence about their effect on the level of productivity within companies (Brynjolfsson & Hitt, 1996; Brynjolfsson & Hitt, 1998; David, 2016; Stratopoulus, 2000). Nonetheless, managers and CEOs have been keeping upgrading and expanding the IT capabilities of the companies they work in, particularly in the field of business analytics.
More recently, advances in the fields of Artificial Intelligence and a greater capacity to store and analyse big amounts of data have together created the need for new tools and new techniques to make sense of the trillions of data points that may be stored in a database. This has, in turn, given rise to the Big Data movement (McAfee & Brynjolfsson, 2012), which calls for new skills that can harness what many describe as a revolution in the way we manage information (Baesens et al., 2016; Kimble & Milolidakis, 2015; Lohr, 2012; Mayer-Schonberger & Cukier, 2014; McAfee & Brynjolfsson, 2012).
This Big Data approach to business analytics is enjoying widespread adoption, not only because it is recommended by scholars (Mayer-Schonberger & Cukier, 2014; McAfee & Brynjolfsson, 2012; Sanders, 2016), but also because its adoption has been shown to provide benefits for companies and the managers who run them (Migliore, 2017). Achieving better performances often requires organizational (Bloom, Brynjolfsson, et al., 2014; Günther et al., 2017) and behavioural changes (McAfee & Brynjolfsson, 2012), that stem from the different nature of the analysis that can be performed with huge amounts of data (Baesens et al., 2016).
As a matter of fact, data-savvy companies are presently capable of monitoring the habits of customer and clients in great detail, and to perform experiments on them every time they visit the company's webpage. These experiments are more focused on finding a useful correlation rather than the right causation, and not only they require a change of mindset towards a data-oriented approach to business, but also raise ethical questions (Luca, 2014).
Facebook and OkCupid have already shown how easy it is to modify a person's behaviour by tweaking an algorithm (Luca, 2014), exposing the degree to which human attitudes and bias are already influenced by the technologies we use. Indeed, the influence that technology has on society and human behaviour may have roots as old as civilization itself (Selin, 2008; Singh, 2011), and Big Data and AI are no exception.
A fierce debate over the effects that these tools are having inside and outside companies is already taking place (Boyd & Crawford, 2012; Ekbia et al., 2015; Marr, 2017; Migliore, 2017; Shah, Horne, & Capellá, 2012). Among those who welcome these new technologies and focus on the benefits they bring, some have uncovered the features of human behaviour that can

be improved as a consequence of incorporating more data and data-focused tools into daily mansions (McAfee & Brynjolfsson, 2012; Migliore, 2017).

Concurrently, human behavioural and organizational methods can be modified to derive the maximum potential value from these technological developments, be that value higher productivity, decreased cost, higher returns, or better analytical insight. Harvard Business Review writers populate the pages of the magazine with suggestions on how to achieve such objectives in several issues (McAfee & Brynjolfsson, 2012; Wessel, 2016).

As a matter of fact, although it has been reported that leaders are aware of these new technologies, there are few peer-reviewed articles that deal with the leadership dimensions of decision-making required to achieve success with Big Data and Machine Learning practices (Migliore & Chinta, 2016).

Indeed, despite this disruption at both company and individual level, research is still scarce on the effects that recent types of AI can have on the human mind, particularly in regard to the decision-making process. Research mostly focuses on the dangers that these technologies pose (Newell & Marabelli, 2015; Boyd & Crawford, 2012) and the issues that may arise when companies integrate them into their business (Baesens et al., 2016; Gopalkrishnan et al., 2012).

Consequently, it is still unclear whether the decision-making process of an individual in a company is affected by the introduction of advanced AI and Big Data technologies, and if it is, what kind of changes it undergoes.

This is nonetheless a critical issue, as without knowing how decisions are taken companies may be exposed to biases, wrongful use of information, and harmful choices. Moreover, these decisions may not only hurt companies' returns but also society, since they result in a misallocation of critical resources (Baesens et al., 2016).

Furthermore, as the subjects of this research are students, the results of such a study could also provide indications about the effects university courses on Big Data and AI have on students' mindsets and logic.

There is already some research dealing with the effects that courses which curriculum includes data analysis and management can have on the people who attend them, although existing studies focus mainly on data literacy (Dunlap & Piro, 2016; Kippers et al., 2018), and do not look at other tendencies that could possibly be affected.

This approach is quite limited, and it is consequently important to explore other tendencies that decision-makers exhibit when making choices, along with the way they are affected by new technologies.

The choice of the tendencies to be inspected, however, is not so straightforward, given the number of possibilities. Just by considering cognitive biases, for instance, more than 186

different tendencies could potentially be explored (Wikipedia, 2018).
Some works have indeed dealt with similar topics, such as the changes in the way data is used in the wake of Big Data (Brynjolfsson & McElheran, 2015). Erik Brynjolfsson, in particular, has produced much work on the effects that new technologies have on individuals, companies, and society (Brynjolfsson & McAfee, 2014; Bresnahan, Brynjolfsson, & Hitt, 2002; Tambe, Hitt, & Brynjolfsson, 2012; Brynjolfsson & McElheran, 2016).
Notably, in the article *Big Data:The Management Revolution*, he and McAfee describe the behavioural and decision-making characteristics that a manager ought to possess in the age of Big Data to improve company performance and increase creativity along with innovativeness.
These features can then be considered the optimal tendencies that individuals and organizations should strive for and are therefore an appropriate choice for the present analysis, as they present interesting learning opportunities.
A confirmation that these tendencies are actually developed in the context of Big Data and AI would signify that individuals are up to the challenges posed by these new technologies, and that companies can then expect to extract high value from them.
Moreover, from an academical point of view, it would imply that the courses offered by the institutions the students belong to are adequately preparing them for the jobs of the future.

Brynjolfsson and McAfee identify various optimal characteristics, of which four will be the focus of this thesis.
First of all, they suggest the deprecation of the old habit of taking decisions with a HiPPO mindset. As a matter of fact, when data are scarce, usually it is the Highest Paid Person's Opinion that matters and dictates the final decision. With Big Data and the means to consume them, consistently taking decisions in this way can be dangerous.
Moreover, they advise companies to start asking "What do we know?" instead of "What do we think?". This means looking at the data first, instead of developing an opinion and only then gathering the data to confirm what is already believed to be true. This kind of confirmation bias is widespread among executives, who often take decisions with a HiPPO approach and then find the data to justify them. Migliore (2017) adds to this topic by advising leaders to continuously test their assumptions, to avoid falling in this kind of bias.
These concepts can be considered as part of an individual's degree of data literacy, as they concern the ability to take correct decisions on the basis of the available data.
Another behaviour that is seen as critical for the successful exploitation of a vast quantity of information is the minimization of the Not-Invented-Here syndrome, which can hinder cross-functional cooperation and the free flow of ideas.

The 4th change that this thesis deals with has not yet been formally identified, but it is seen as a possible development of Brynjolfsson and Hitt's paper *Information Technology and Internal Firm Organization: An Exploratory Analysis* (1997). In the paper, the authors show how investment and greater adoption of IT in firms increased the degree of decentralization and authority delegation in American firms.
With the advent of ever more intelligent software, it could be possible for the extent to which a manager is willing to delegate decisions to a machine to increase, especially as computers are seen as the new subordinates, capable of executing complex tasks.
A more detailed introduction to the phenomena studied in the present thesis follows.

2.5 Phenomena Studied in the Present Research

2.5.1 Data Literacy

Data literacy is a recent evolution of the concepts of statistical literacy and information literacy (Schield, 2004), from which it inherits most of its characteristics.
The elements that comprise it are differently defined in different studies (Dunlap & Piro, 2016; Qin & D'Ignazio, 2010; Kippers et al., 2018), but at its core it concerns the ability to correctly analyse, interpret, and evaluate information and data, a skill necessary to adequately perform data analysis (Schield, 2004).
The importance of possessing data literacy is hard to understate, as companies have now access to vast and ever-increasing amounts of data (Mayer-Schonberger & Cukier, 2014) and are in need of expertise to harness their potential benefits. This need is further emphasized by the recently demonstrated benefits of data-driven decision-making (Brynjolfsson & McElheran, 2015; Brynjolfsson et al., 2011), which calls for data literate employees and managers who can take data-based decisions to unlock new value for their company.
It therefore comes as no surprise that previous works have tested different aspects of data literacy and the way it can be changed. For instance, exposure to data and the ways they can be handled has been found to affect people's use of data (Dunlap & Piro, 2016; Kippers et al., 2018), increasing their ability to reach accurate conclusions after a fitting analysis of the available information.
In this thesis, data literacy is analysed through the testing of the respondents' tendency to rely on data to take correct decisions during their decision-making process, and is tested together with confirmatory bias and HiPPO, as the latter are factors that could hinder it.

2.5.2 Confirmatory Bias

Confirmatory Bias is a particular kind of cognitive bias, a group of patterns that humans, and also some animals, can follow systematically in an attempt to reach a decision or form a judgement (Haselton, Nettle, & Andrews, 2015). These patterns, adopted to save mental energy, usually deviate from rationality, and affect the final judgement in often a negative way.

In particular, the person that experiences a confirmatory bias is more prone to search and accept information that confirms previously held beliefs or expectations (Jonas, Schulz-Hardt, Frey, & Thelen, 2001). Moreover, confirmatory bias leads to the misinterpreting of ambiguous evidence so to confirm the person's current hypothesis about the world, be it a prejudice or an expectation (Rabin & Schrag, 1999).

Another characteristic of the confirmatory bias, called the primacy effect, is that when a person must draw a conclusion on the basis of information acquired and integrated over time, the information acquired early in the process is likely to carry more weight than that acquired later (Nickerson, 1998).

While in the present study the information given to the research subjects is introduced over a short period of time, as required by the survey design, it is nonetheless a key element of the bias that is manifested in real-world situations, where information is more likely to be acquired over a longer period of time.

This process, called sequential information seeking, has been understudied in scholarly literature, which experiments have mostly utilized simultaneous information search as the method of feeding information into the people under examination (Jonas et al., 2001). While this does not reflect how people outside the laboratory process most information (Jonas et al., 2001), it is a design necessary for the time and budget constraints that most studies are subject to.

Confirmatory bias is regarded as a serious problem mainly because a person can reach a conclusion or belief that does not reflect reality and will keep acting considering said belief to be true. Even when this does not occur, a distorted vision of reality can be dangerous as it leads to a flawed decision, especially when a decision has financial consequences. This issue can manifest itself from the very early stages of information search and can lead a person to overlook risks and warning signals even when they are evident (Jonas et al., 2001). People will end up seeing patterns in the data regardless of whether the patterns are really there (Nickerson, 1998).

While this is a problem often overlooked in business, the dangers of confirmation bias has been well documented by research (Nickerson, 1998; Schulz-Hardt, Frey, Lüthgens, & Moscovici, 2000), which also found out that confirmatory bias does not disappear when the subject who needs to process information and take decisions is a group of people rather than a single individual (Schulz-Hardt et al., 2000).

For simplicity, in the present study only one actor taking a decision is analysed.
The early information fed to her, that should induce the bias, is that her manager, a highly competent and respected person in the field, has a certain opinion about a decision that has to be taken. The data that is then presented to the respondent can either confirm or deny the earlier information. If a confirmatory bias does indeed form, the survey design should make it possible to measure its intensity.

2.5.3 Highest Paid Person's Opinion

Traditionally, decision-making in companies has been the domain of experts rather than data. This happened because data were traditionally scarce, and people that had developed useful heuristics and decision patterns from experience were particularly appreciated and rewarded (McAfee & Brynjolfsson, 2012).
However, the new data avalanche brought by the Big Data movement calls for new tools and processes that can take advantage of the surge in data availability (McAfee & Brynjolfsson, 2012).
Although anecdotes about HiPPOs abound, few authors explore the phenomenon in detail, and most of the time references to them can be found not on peer-reviewed journals, but on articles meant for an audience of managers and business executives (Lohr, 2012; McAfee & Brynjolfsson, 2012). This probably hints at the very practical nature of a HiPPO, a business routine or practice that has not been properly formalized by the literature.
The authors that deal with HiPPOs, mainly in the process of advising managers on best business practices, stress how they can be a problem in the age of data-driven decision making and advanced analytical tools. Intuition and experience, although useful, are not the most appropriate tools to use when vast quantities of data to explore options and perform analysis are available, and the Highest Paid Person Opinion, or HiPPO, may prevent a company from gaining full value from its data (Lohr, 2012; McAfee & Brynjolfsson, 2012).
To avoid this issue, students should be taught to critically consider their superiors' opinions, so that HiPPOs are not given more relevance than necessary during data-driven decisions.
Thus, one the goals of this thesis is to verify whether universities can effectively prepare

students to correctly assess the importance and validity of an executive's opinion when data are available.

2.5.4 Not-Invented-Here Syndrome

Katz and Allen (1982) define the Not-Invented-Here (NIH) syndrome as the tendency of a group of stable composition to believe it possesses a monopoly of knowledge in its field, which leads it to rejecting ideas from outsiders to the likely detriment of its performance.

More generally, the NIH syndrome can manifest itself when a group working on a project comes into contact with external knowledge, technology or solutions for the task at hand. When confronted with external technology, a group is forced to compare its expertise to the one of the external source of technology. This process can constitute a threat for the group's perceived expertise, and for the integrity of the self-image that group has of itself (Wastyn & Hussinger, 2011).

As a consequence, the group can become resistant towards the incorporation of external technology into its project. This is particularly relevant when groups are composed of people who identify strongly with the group itself, since a confrontation with the external environment can evoke feelings of inferiority that are felt deeper and more mightily threatens the self-esteem of the group (Wastyn & Hussinger, 2011).

The same happens at an individual, rather than group, level (Agrawal et al., 2010; Piller & Antons, 2015), as NIH Syndrome is a psychological, not only social, sentiment of hostility towards external knowledge (Husted & Michailova, 2002).

This phenomenon was first observed and became widely known in the R&D community of engineers (Katz & Allen, 1982), before it was reported in the academic literature by Clagett (Clagett, 1967), who had spent time in an engineering research centre (Wastyn & Hussinger, 2011).

Starting from that insight, the topic has been studied more in depth, particularly because of the detriment that it can cause to innovativeness and the capacity of companies to absorb external knowledge. However, most of the research has focused on anecdotal aspects of the phenomenon, rather than describing it in a more scholarly fashion (Piller & Antons, 2015).

One of the most important contributions is that of Katz and Allen (1982), who expose how this kind of behaviour can be dangerous for business. Professionals and organizations need to interact with external sources of technology, inspirations and ideas if they want to maintain technological adequacy, if not technological edge. NIH syndrome discourages them from doing so, endangering their work.

The authors also show how NIH Syndrome does, as a matter of fact, reflect in project performance. The latter tops when a group has not had a long tenure together, and the longer the tenure of a group, the worse the project performance. Indeed, the members of a group that is maintained for a longer time increase their self-identification with the group itself, thus increasing the likelihood that the group will resist external technology. This prevents the project from benefiting from external ideas, making it more prone to failure, and confirming NIH syndrome's nefarious effects.

These effects are more visible when the external suggestion comes from similar groups, like competitors, while they hardly emerge when the external agents proposing the technology are customers, suppliers or universities. The grade of similarity seems indeed to directly affect the emergence of NIH Syndrome (Wastyn & Hussinger, 2011)

That is the reason why in the scenarios about NIH syndrome used in this research the source of knowledge is an external firm that makes similar products. This is advisable in order to make NIH syndrome more likely to emerge, so to make it possible to estimate the change in its levels across time and treatment. For the same reason, the firm the respondent belongs to in the scenario is successful, as this also increases the likelihood to exhibit NIH syndrome (Wastyn & Hussinger, 2011).

If a bias for preferring internal over external innovation is indeed present in the students and reduced by the treatment, the aforementioned elements should make its status more easily detectable.

2.5.5 Delegation

Delegation within companies is a topic that has been present for some time in the scholarly literature, especially in firm organization theory (Hitt & Brynjolfsson, 1997).

Colombo and Delmastro (2004) identify 3 families of factors that can influence the allocation of decision-making power between a manager and his superior.

Among all of these factors, the flow of information is perhaps the one which is most carefully studied. The topic has undergone extensive investigation due to the major developments in information communication technology (ICT) that have changed the way information is broadcast and exchanged within a firm (Bloom, Garicano, Sadun, & Van Reenen, 2014).

In particular, ICT makes it possible for information and decisions to flow rapidly within the company, decreasing the need for centralization to increase the velocity and accuracy of collaboration among employees. As a matter of fact, increasing decentralization in an environment where adequate ICT tools are available can improve knowledge sharing, since highly specific knowledge is more likely to reside at the lower levels of an organization (Hitt & Brynjolfsson, 1997).

Firms seem to be well aware of these advantages, as various research has found out that the level of investment in IT strongly correlates with the degree of delegation and decentralization within the company (Colombo & Delmastro, 2004; Hitt & Brynjolfsson, 1997; McElheran, 2014). When information is more easily accessible, managers tend to delegate decisions to subordinates, who are the ones that, at the local level, have the knowledge required to solve the task.

What happens, then, when the entity with the knowledge required to solve an issue or perform a job is not a person, but a machine? Humans are limited as information processors (Hitt & Brynjolfsson, 1997), and in a world where everything is becoming quantifiable and the amount of data that can be processed is increasing by orders of magnitude (Mayer-Schonberger & Cukier, 2014; Beer, 2016), machines are the best suited for processing information.

Computers, however, do not only process information in order to feed it to a human controller, but also take autonomous decisions (Newell & Marabelli, 2015), like deciding whether to negate a loan to a person based on her credit score.

These new capabilities are causing a shift in the traditional view about computers, increasingly considered less as tools and more as autonomous agents (Hoc, 2000), thus more easily comparable to a human employee.

If software and algorithms are able to take decisions or provide invaluable information to a human to do so, then human managers could delegate tasks related to decision-making to their computers. If this is the case, a tendency parallel to the one described by Hitt and Brynjolfsson (1997) could be noted. That is, a higher availability of advanced analytics technologies, built on Big Data and Machine Learning, should increase the degree of delegation towards machines. As humans are able to communicate in a better way with machines, and machines possess more specialized knowledge and data, delegation seems the most appropriate way to meet the challenges posed by Big Data.

This particular topic has been scarcely studied by academia, as the technologies that make it possible to delegate complex decisions to machines are very recent, and most companies do not have the expertise required to use them.

Therefore, the present research also aims to expand the knowledge about the effects that being exposed to Artificial Intelligence has on the propensity of managers, or people in a position to make a decision, to delegate part or all of a decision-making task to a software.

3 Method

3.1 Research Design

The main purpose of this thesis is to uncover whether the exposure to tools and methods concerning AI and Big Data analytics can change the way people in a position of power take decisions.

To do so, the experimental design took inspiration from the difference-in-differences method. A difference-in-differences method is an experimental design commonly used in medical research for the trial of a new medicine, or in economics for policy analysis (Bertrand, Esther, & Sendhil, 2004).

The subjects are divided into 2 groups, usually without them knowing which group they belong to, and a treatment (like a new medicine) is applied to only one group. Before applying the treatment the characteristics of both groups are measured, in order to provide a measure of their current values in the dimensions of interest. These could be for instance their red blood cells count, or the spread of an infection within their bodies.

Later, some time after the treatment has been applied, the groups are tested again to see what, if anything, has changed in the values of the dimensions of interests. Then, the difference in the values between the second and the first measurements of the group that did not receive the treatment are compared to the difference in the values between the second and first measurements of the group that did receive the treatment. This way, it is possible to isolate the changes that are due to the treatment itself from the global changes that may have affected both groups.

One of the main advantages of this method is that it allows for the minimization of the self-selection bias, as people with certain innate features that could affect the results could be more likely to take part in the experiment. This happens especially because ethics require experiments performed on people to have voluntary participation, and unless the participants are tested beforehand, it is hard to find out whether they are a sample representative of the population or not.

This method was therefore chosen for a variety of reasons:

- Since the measurements are taken before and after the treatment is applied, it enables the tracking of how the decision-making process changes, if it does, in time. In this case it is particularly relevant since the treatment was applied during the course of an entire semester.

- It diminishes the possibility of a pre-existing selection bias. Since the treatment, that is to say attending a course on AI or Big Data, is not randomly allocated, but students can choose whether or not taking this kind of course, this is particularly relevant in the present research.

- With this design there is no need to pre-emptively vet the participants to exclude those who possess different features from the others that could drastically change the results. Since the potential sample size is not large to begin with, this admits more survey respondents and therefore more significant findings.

3.1.1 Sample Selection

While testing managers may have been more relevant for the research question, applying a treatment that can expose the subjects to in-depth knowledge of AI and Big Data would have been hard.
Therefore, business school students enrolled in a master's degree were chosen as a substitute, a practice which is not new to research (Hutchinson et al., 2010; Schulz-Hardt et al., 2000; Aiman-Smith, Scullen, & Barr, 2002).
In this study, the students considered were all enrolled in either Bocconi University or Copenhagen Business School.
The reasons for the choice of this particular sample are various:

- Although not the managers of today, business students are likely to be the managers of tomorrow, or at least to resemble business managers because of their study choice.

- The current structure of the EMIT (Economics and Management of Innovation and Technology) and MIB (Management of Innovation and Business Development) masters curricula allows a "natural experiment" to occur in the use of the treatment. Students can choose to enrol in courses that are part of a specialization in Big Data, which also includes advanced analytics and AI among its topics. Therefore, some of them are exposed to these tools at a university level, which is assumed to be more than amateurial and can be considered likely to have an effect on the decision-making process of the students, if any effects are indeed present.

- As the respondents come from very similar academic backgrounds, they can be seen as a relatively homogenous population. Therefore, the effects of a treatment should be more likely to arise consistently across the test subjects.

- As a population, they are more approachable with the means possessed by the author and can therefore provide a higher rate of response than, for example, company executives.

The most important reason for the choice of this population is number 2, here explored in more detail.
Bocconi and CBS offer similar programs, as testified by the number of double-degree agreements that the two universities have. In particular, they both offer a specialization that is focused on Big Data and advanced analytics. The Major in Big Data and Business Analytics at Bocconi and the Minor in Data in Business at CBS offer courses where students are exposed to advanced analytical tools that are part of the technologies described by Brynjolfsson and Migliore (McAfee & Brynjolfsson, 2012; Migliore, 2017).
Taking this specialization is therefore considered an appropriate treatment for this experiment. A potential drawback is the fact that the treatment cannot be randomly allocated, as students are free to choose whatever courses they prefer. To reduce the self-selection bias that may arise as a result, the modified difference-in-differences method is employed.
Therefore, students were administered 2 surveys, the first before the application of the treatment and the second after, the former to measure their initial decision-making tendencies and the latter to measure how they had changed after the treatment. The first survey was administered during August and September 2017, while the second survey was administered during February and March 2018, after the end of the semester during which the treatment courses were held.
There was no initial formal group division, but the students who answered the second study were divided into the treatment and the control group based on whether they had taken courses about AI, Big Data, or advanced business analytics. The structure of the survey made it possible to track the answers of the individual students, so to assign them to the treatment and control group retroactively.

3.2 Survey Structure

The study was conducted by administering 2 surveys with an identical structure, aimed at measuring the decision-making tendencies of the respondents along 3 dimensions before and after the treatment.
The surveys were built using the Qualtrics platform, and distributed via link through instant messaging services, emails and Facebook groups of business students.

Only the students who successfully completed the first survey were given the second questionnaire. The first survey was answered by 67 students, while the second collected 50 responses, for a total of 117 observations.
As the topic of analysis is quite complex, a vignette design was adopted, the most suitable for this activity (Aguinis & Bradley, 2014). A brief description of what this kind of design entails follows.

3.2.1 Experimental Vignette Methodology

Experimental Vignette Methodology (EVM) is a type of survey design that has been used since at least the 1970s (Finch, 1987). Respondents are presented with a "vignette", usually a written statement describing a situation in which the respondent could find herself, preceding a question about the beliefs, feelings, or decisions that would be experienced in that situation by the respondent.
The response to the question is usually the dependent variable that is being measured, while the vignette structure can include different independent variables that can be manipulated by the researcher in order to obtain different vignettes. These different vignettes facilitate the inference of causal relationships between the independent and dependent variables.
This design has been used mainly in sociology, management and organizational behaviour studies (Finch, 1987; Rooks, Raub, Selten, & Tazelaar, 2000), although its use in medicine and psychology is not uncommon (Ross, Moffat, McConnachie, Gordon, & Wilson, 1999; Thompson, 2003).
The fact that the respondent is presented a situation where she can immerse herself strikes a balance between field studies, where the environment could influence the results beyond understanding, and laboratory experiment, where the strict design may leave out important explanatory variables (Aguinis & Bradley, 2014). This characteristic makes the method fit for measuring implicit beliefs and tendencies that, if asked explicitly, the participants would be reticent to answer honestly, such as discrimination or business ethics.
Another benefit of EVM is that it allows researchers to understand the decision-making process of a single individual, and thus allows a study to gather useful amount of information even from a small population or sample (Aguinis & Bradley, 2014).

Different designs of EVM are possible, and the literature identifies 3 types in particular: between-person, within-person, and mixed (Atzmüller & Steiner, 2010).
In between-person designs each participant reads the same vignette, and the answers of each

participant are then compared. Participants in within-person designs also read the same set of vignettes, but the comparison is then made between vignettes within the same respondent (Atzmüller & Steiner, 2010).

Mixed designs, on the other hand, divide the respondents in different groups to which different set of vignettes are showed, and comparisons are then made across the different groups. Since multiple respondents offer responses about the same vignettes, they can be compared (Aguinis & Bradley, 2014).

This research adopts a mixed design, as participants are randomly divided in 4 groups, according to the particular set of manipulation within the vignette they are viewing, and their answers are then compared to identify their tendencies in regard to the dependent variables, both before and after the treatment.

Although it is possible to create vignettes out of videos and images, or even Virtual Reality (Aguinis & Bradley, 2014), the present study opts for a more traditional text-based approach for 2 main reasons. First, it was the easiest and cheapest to implement. Second, this approach lets the respondent use more of his imagination to be inserted in a setting that, although familiar, may never have been personally experienced.

The main reason for using a vignette design is its usefulness in measuring the respondents' reliance on a certain factor during their decision-making process.

Each scenario is defined by two variables, which can assume two values. If a respondent does not consider an element to be important for his decision, then her answer should not be affected by a change in the value of that variable. On the contrary, if a respondent bases her answer on a certain element, then a change in the value of that element should considerably change the answer as well.

Thus, by comparing how much a change in each element contributes to a change in the final answer, it is possible to measure the relative importance and influence of an element on the decision-making process of the respondent.

However, this may not be sufficient to test *Hypotheses 2* and *3*, as the tendencies that need to be examined for their investigation would be measured only indirectly. As a matter of fact, the strength of this design may be more apparent in the measurement of the respondents' reliance on the factors considered, rather than in the detection of the respondent's affinity for said factors.

Therefore, the results from the regression models were also integrated with the results obtained from an analysis of the means of the values assumed by the respondents' answers in the different versions of the same scenario. These additional statistics grant a direct measurement of the students' preference for one of the two options proposed in each scenario, as they indicate the likelihood of a respondent choosing one option rather than the other.

For the aims of this study, this kind of design provides additional benefits:

- It is possible to articulate in detail the different situations that can elicit a decision-making process in a manager, while at the same time controlling the variables that form the situation whence the respondent is required to take a decision.
- Scenario descriptions provide a way for business students to immerse themselves more deeply and realistically in real-life situations. As a matter of fact, they rarely have first-hand experience of instances in which they would have to make a critical decision for a company. Scenarios, on the other hand, allow students to put themselves in a state of mind that resembles that of managers.
- Manipulating the variables in the vignettes, using a 2x2 format, allows a more detailed and controlled setup of the scenarios. As mentioned before, this eases the emergence and measurement of the implicit tendencies in the decision-making process of the respondents by confronting the different answers given to different variables' combinations.

3.2.2 Vignettes and Clues Used in This Research

Each scenario contains 2 independent variables that can assume 2 different values, high and low or weak and strong, thus allowing for the creation of 4 different vignettes per scenario. The surveys contain 3 scenarios, one for each of the tendencies that are being measured. The total number of different vignettes per survey is therefore 12.
Across the studies conducted with EVM, this number may fall on the low end, as studies with hundreds of different vignettes have been published (Aguinis & Bradley, 2014). However, more vignettes would have increased the risk of information overload and fatigue for the respondent (Weber, 1992), especially considering that the descriptions of the scenarios can be quite long.
While the particular setting and writing of a scenario for a given dimension differs from one survey to the other, they are nonetheless built so as to measure the same tendency and using similar variables. The difference in the setting of the scenarios between the 2 surveys is needed in order to prevent students from recognising what is being measured, and also to preclude them from remembering the previous answers they gave. Had it happened, they could have been biased during the completion of the second survey.

In both surveys, each respondent views only one version (or vignette) of a given scenario. The vignette assignment is randomized, so it is possible for the same respondent to answer different versions of the same scenario in the two different surveys. This provided the means for a more significant analysis of the mean answers.
Each vignette contains a description of the scenario and a question. Since the question that is being asked is a binary choice between two options, but a more detailed degree of choice is advisable, a draggable bar, or visual analog scale, is employed.
A visual analog scale (VAS) is preferred to a Likert Scale mainly because the results were to be analysed using regression analysis and Pearson correlations. These techniques are not well-suited for responses obtained with a Likert Scale, which can also easily cause problems when handling ordinal data (Bishop & Herron, 2015).
The bar goes from 0 to 100, and the initial position of the marker is in the exact middle of the bar, with a value of 50. The two different options towards which the respondent may lean are displayed at the two extremities of the bar.
The respondent can then drag the bar towards one of the two extremities, thus expressing her leaning towards one of the two options with a level of granularity much higher than in a Likert Scale.
An example of such setting can be seen in the following image:

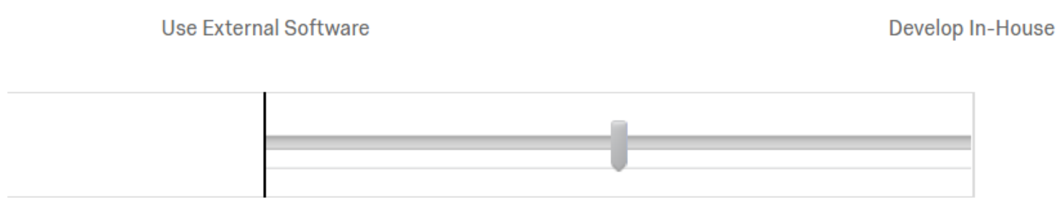

The scenarios and the variables, or clues, that are manipulated to build different versions of the same scenario, are the following.

Scenario 1 The decision-making process measured in the first scenario is the tendency of relying on HiPPO and confirmatory biases rather than data, or vice-versa.

The respondent first receives the opinion of a superior about the success of a product and is then asked to evaluate the available data to suggest a course of action. In this case, the HiPPO constitutes the first informative piece of information received and is the element that should form the confirmatory bias in the mind of the respondent. As the HiPPO is part of the confirmatory bias, and both phenomena should bring the respondent to take a similar course of action, they have been grouped together for analysis in this research.

The data were intentionally given a certain degree of ambiguity, as confirmatory bias is more likely to appear when a person has to take a decision based on ambiguous evidence (Rabin & Schrag, 1999). Therefore, showing data that are too obviously leaning towards one option or the other would eliminate the need for a bias to be used. As a matter of fact, one would need not to utilize a confirmatory bias when considering very obvious data, as the decision would require a small amount of mental power and would be reached quickly without employing bias or heuristics (Benson, 2016).
If a bias is indeed present, a respondent's tendency to adjust her decision towards the opinion of the senior executive, even when the data say otherwise, should be observed.

The clues, or independent variables, are:

1. Level of confidence of the senior executive in the success of the product. It can assume the values High or Low.

2. What kind of market the data are showing. It can assume the values Favourable or Not Favourable. This clue can also be considered as the likelihood of success the product actually has.

The possible combinations are the following:

<table>
<tr><td></td><td></td><td colspan="2">Confidence of Senior Executive in Success</td></tr>
<tr><td></td><td></td><td>High</td><td>Low</td></tr>
<tr><td rowspan="2">Actual Likelihood of Product Success</td><td>High</td><td>High Likelihood of Success
High Confidence in Success</td><td>High Likelihood of Success
Low Confidence in Success</td></tr>
<tr><td>Low</td><td>Low Likelihood of Success
High Confidence in Success</td><td>Low likelihood of Success
Low Confidence in Success</td></tr>
</table>

Table 1: Scenario 1 cues

The question asked, or the dependent variable, is the willingness of the respondent to suggest the product launch to its superior. The respondent is allowed to choose where to position the

cursor on the draggable bar between two extremes, expressing his likelihood of suggesting either of the proposals.
The left end of the bar conveys the option *Unlikely to suggest the launch of the product*, with a value of 0. The right end conveys instead the option *Likely to suggest the launch of the product*, with a value of 100.

Scenario 2 The decision-making process measured in the second scenario is the Not-Invented-Here Syndrome, or the tendency to rely on in-house rather than external innovation.

The aim of this scenario is to measure the tendency of a respondent to rely on in-house innovations rather than on external innovations, even when the latter can bring more benefits. In the vignettes the respondents are provided with an internal and an external technology to solve the problem they are facing, namely the development of a new product. Each technology can be evaluated not only on its fit to the current product, but also on its feasibility and convenience for the business at large.
Therefore, the mere potential technical preference of the respondent should not influence the answer, as the survey requires a more comprehensive view of the outcomes that choosing each technology will bring.
In order to make the NIH syndrome more likely to emerge so to measure it, the scenario has been constructed according to certain criteria that are more likely to elicit it (Wastyn & Hussinger, 2011).

The criteria are:

1. The firm the respondent belongs to in the scenario is successful.

2. The source of the external technology is a firm similar to the one the respondent belongs to, but not a customer, supplier or university.

The clues, or independent variables, are:

1. Strength of external opportunity. It can assume the values Strong or Weak

2. Strength of in-house opportunity. It can assume the values Strong or Weak

The possible combinations are the following:

		Strength of external opportunity	
		Strong	Weak
Strength of in-house opportunity	Strong	Strong in-house opportunity Strong external opportunity	Strong in-house opportunity Weak external opportunity
	Weak	Weak in-house opportunity Strong external opportunity	Weak in-house opportunity Weak external opportunity

Table 2: Scenario 2 cues

The question asked, or the dependent variable, is the preference of the respondent for either the external or the internal innovation. The respondent is allowed to choose where to position the cursor on the draggable bar between two extremes, expressing his likelihood of suggesting either of the proposals.
The left end of the bar conveys the option *Use technology developed externally*, with a value of 0. The right end conveys instead the option *Use technology developed in-house*, with a value of 100.

Scenario 3 The decision-making process measured in the third scenario is the tendency to delegate tasks to machines rather than doing it in person.

In this scenario, the machine is represented by a software with varying degrees of capability and appropriateness for the analysis that has to be performed.
The human, the other actor that could accomplish the task exposed in the scenario, is also described in each vignette with varying degrees of skills. The respondent should therefore be able to identify himself with the person described in the vignette and his skills and take a decision accordingly.

The clues, or independent variables, are:

1. Strength of human (self) skills for solving the issue. It can assume the values Strong or Weak.

2. Strength of machine skills for solving the issue. It can assume the values Strong or Weak.

The possible combinations are the following:

		Strength of human(self)	
		Strong	Weak
Strength of machine	Strong	Strong machine Strong human	Strong machine Weak human
	Weak	Weak machine Strong human	Weak machine Weak human

Table 3: Scenario 3 cues

The question asked, or the dependent variable, is the preference of the respondent for either solving the issue himself or delegating the task to a software. The respondent is allowed to choose where to position the cursor on the draggable bar between two extremes, expressing his likelihood of suggesting either of the proposals.
The left end of the bar conveys the option *Do not delegate the task*, with a value of 0. The right end conveys instead the option *Delegate the task to a software*, with a value of 100.

4 Analysis

4.1 Methods of Analysis

The analysis of the results was conducted with SAS 9.4, which provided the tools for performing regression, correlation and variation analysis of the respondents' answers and characteristics.
The cues of each scenario shown to the respondents are represented as binary variables. In particular, if an attribute assumes the value *high* or *strong* in a scenario, it is represented by a 1, while if it assumes the value *low* or *weak* it is represented by a 0. This later allowed the development of regression models based on dummy variables.
The answers that the respondents gave to each scenario are standardized for the regression analysis, to achieve a better representation of the coefficients and other measures obtained from it.
A standard OLS multiple linear regression is used first on the data of the first survey, and then on the data of the second survey. The differences between the results of these analyses are used to discuss the validity of the hypotheses of this research.
As the cues in each scenario are randomized, so that each respondent has the same probability to receive a 0 or a 1 for each of the two independent variables, the risk of multicollinearity is expected to be rather low.
In the regression analysis each scenario is inspected separately, and in the construction of the model respondents are considered according to their belonging to the treatment or to the control group. This enables a more detailed investigations of the phenomena and tendencies considered.
The regression model is built as follows:

$$y_{ij}^{(k)} = a_{1j}^{(k)} c_{1ij}^{(k)} + a_{2j}^{(k)} c_{2ij}^{(k)} + b_{ij}^{(k)} + \epsilon_{ij}^{(k)} \tag{1}$$

Where y is the respondent's answer, a_1 is the first coefficient, c_1 is the first clue, a_2 is the second coefficient, c_2 is the second clue, b is the intercept, ϵ is the error, i is the respondent, j is the scenario and k is the group considered. c_1 and c_2 are categorical variables that can assume a value of either 0 or 1.

In addition to the regression models, the mean responses in each version of each scenario

were compared to highlight possible differences in their behaviour between control and treatment group.
The mean values of the responses in the first survey were subtracted from the mean values of the responses in the second survey, to analyse the change these means went through, if any was to be found.
This was done separately for each version of each scenario and for each group. As there are 4 versions for any given scenario, and 2 different groups the respondents can belong to, each scenario produced 8 different values.
Thus, it was possible to obtain the change the mean value of the responses went through in the course of the six months between the surveys, for any group and cue combination.

5 Results

During the exposition of the results, the means and standard deviations of the values assumed by the respondents' answers in the different scenarios are shown together with the results of the correlation analysis, as the tables provide an overall view of the characteristics of the scenarios' observations.
The overall high standard deviations are caused by the the fact that each scenario could be presented to the respondent in 4 different versions, which could result in very different answers, as already exposed in the section on the method. To reduce this uncertainty, the means of each version of each scenario are separately exposed and analysed in the sections of the results which focus on separate scenarios.
Moreover, an analysis of the medians confirmed the tendencies and behaviours that emerge from the means, so that showcasing and discussing the former would be superfluous.

5.1 Results of Correlation Analysis

Control Group								
	Before				After			
Measure	Mean	Std.Dev	1	2	Mean	Std.Dev	1	2
1. Scenario 1	63.590	21.327			62.500	23.553		
2. Scenario 2	53.205	25.935	0.1535		43.500	26.567	-0.0565	
3. Scenario 3	59.590	28.770	0.3963**	-0.0751	57.077	28.719	0.2774	-0.0689

*p < 0.1, **p < 0.05, ***p < 0.001, ****p < 0.0001

Table 4: Correlations Between Scenario Answers (Control Group)

Considering only the control group, as showed by table 4, we can observe an initial correlation between answer to scenario 1 and scenario 3 (0.39627, p-value: 0.0125), weakened in the second survey, where the correlation coefficient is equal to 0.27740 and its p-value to 0.1701. In general, it seems that while some interesting correlations can be observed in the first survey, after six months all of them have weakened.
In the first scenario the mean of the answers is always higher than 60, and it experiences only a small decrease after six months, showing a general preference in the control group for launching the product.
In the second scenario the mean, initially with a value of 53.205, decreases by almost 10

points in the second survey, suggesting that the respondents in the control group changed their preference towards *External innovation* after six months.
In the third scenario, the means behave like in the first. In both surveys their value hovers quite above 50, and they experience a small decrease from the first to the second survey. This highlights an unwavering preference for delegating the task to a machine.

Treatment Group								
	Before				After			
Measure	Mean	Std.Dev	1	2	Mean	Std.Dev	1	2
1. Scenario 1	64.593	25.341			58.292	26.201		
2. Scenario 2	39.556	28.534	0.0789		40.625	25.483	0.0299	
3. Scenario 3	54.259	35.125	0.0466	-0.1601	58.583	29.469	0.3868*	-0.3953*

*p < 0.1, **p < 0.05, ***p < 0.001, ****p < 0.0001

Table 5: Correlations Between Scenario Answers (Treatment Group)

Considering only the treatment group, no significant initial correlation between answer to scenario 1 and scenario 3 (0.04657, p-value: 0.8176) can be found. This correlation, however, becomes significant and much stronger in the second survey (0.38679, p-value: 0.0619), differently from what happens in the control group.
The same behaviour is showed by the correlation between the answers to scenario 2 and scenario 3 in table 5. The group initially exhibits a negative correlation between the answers to scenario 2 and scenario 3 (-0.16010, p-value: 0.4250), that becomes stronger and much more likely to be significant in the second survey (-0.39531, p-value: 0.0559).
On the other hand, the answers to scenario 1 and 2 show no significant correlation at all.
However, some peculiar features of the treatment group can be observed when looking at the means.
In the second scenario, the mean value is essentially the same in both surveys. This value, however, is much below 50, as it hovers around 40, showing how the treatment group exhibits a distinct preference for the *Use external innovation* option, much more so than the control group.
In the third scenario, differently to what happens in the control group, the mean value increases from the first to the second survey, going from 54.259 to 58.583. This behaviour implies that the treatment group is the only one to increase its preference for the option *Delegate to a software* after six months at the aggregate level, although very slightly.

5.2 Results of Regression and Mean Analysis

5.2.1 Preamble

As the dependent variable was standardized, the values of its observations represent their difference from the mean of the original variable in number of standard deviations. For instance, a standardized value of 0.7 implies that the original value is 0.7 standard deviations above the mean, while a value of –1.3 indicates that the original value is 1.3 standard deviations below the mean (UCLA: Statistical Consulting Group, 2017).
As a consequence, the coefficients obtained as a result of the linear regression show how much the passage from a low to a high value of an independent variable changes the distance of the dependent variable from its mean.
A negative coefficient signifies that, when the independent variable changes from *low* to *high*, respondents increase their preference for the option on the left of the bar.
On the opposite, a positive coefficient indicates that, when the independent variable changes from *low* to *high*, respondents increase their preference for the option on the right of the bar.
Moreover, the intercept is the indicator of the value that a response assumes on average when both independent variables are equal to 0, that is to say when both cues are set to *weak*. This provides an outlook of the theoretically unaffected tendency that the respondents possess. As a matter of fact, if both independent variables are set to the same value, the respondent should be completely or almost neutral in regard to which option to choose, if there are no preexisting preferences.
On the opposite, non-standardized answers were used for the analysis of the means and how they change both in the different versions of the same scenario and across different surveys. The reason for this decision was to provide a clearer view of the means and their values, that are more easily understandable and comparable when not in standardized form, for the purpose of this specific analysis.

5.2.2 Scenario 1

	EXEC = 0		EXEC = 1	
	Control	Treatment	Control	Treatment
DATA = 0	-1.917	2.589	-15.178	-31.095
DATA = 1	0.289	-3.200	-4.708	11.600

Table 6: Scenario 1: Differences in Means (After - Before)

The first notable finding is that the treatment group seems to have acquired, after the treatment, a higher ability to take a decision that reflects the data available, rather than the opinion of the executive.
Although both groups' mean responses do not undergo significant changes across the two surveys in the versions where the executive is not very confident about the success of the product, the same is not true for the versions where the executive is confident about the product's success.
As a matter of fact, the treatment group seems to perform better than the treatment group in recognizing the most appropriate choice for each scenario version in the second survey, taking its decisions in accordance with the data available.

DV : Likelihood to suggest product launch				
	Control Group		Treatment Group	
	Before	After	Before	After
Intercept	-0.4998**	-0.5040	-0.4101	-0.6032**
	(0.2257)	(0.3526)	(0.3676)	(0.2736)
EXEC	0.3578	-0.0026	0.3534	-0.1025
	(0.2815)	(0.34811)	(0.4333)	(0.3511)
DATA	0.6705**	0.8978**	0.4832	1.3478***
	(0.2815)	(0.3659)	(0.4309)	(0.3549)
Adj. R^2	0.1316	0.1399	0.0063	0.3507
F-Value	0.0298	0.0678	0.3548	0.0041
N	39	26	27	24
*p < 0.1, **p < 0.05, ***p < 0.01 ****p < 0.001				

Table 7: Scenario 1 Regression Model: HiPPO, Bias and Data

Looking at the regression results in table 7, it is possible to see how the EXEC variable influence on the dependent variable, strong in the first survey, disappears in the second survey,

as its coefficient turns from positive (0.35774) to negative (-0.00258) and its explanatory power is reduced as its p-value worsens.
The variable DATA is still the most important to explain the variance of the model, and its coefficient increases after six months have passed, going from 0.67052 to 0.89785
Its significance, on the other hand, remains almost identical, as its p-value changes from 0.0226 to 0.0221.
The explanatory power of the model remains however almost unchanged.
Considering only the treatment group, a significant change can be seen in the answers of the respondents from the first to the second survey in regard to their reliance on the different factors used to take a decision.
First of all, the Adj. R^2 of the model increases by more than 30 times, from 0.0063 to 0.3507, despite a smaller sample size. Likewise, the model's p-value goes from 0.3548 to 0.0041. These results show that while in the first survey the variations in the independent variables could not explain the variation in the dependent variable, they can do so in the second survey. Furthermore, the variable DATA, while non-significant in the first survey, becomes a valid explanator of the dependent variable in the second model, as its p-value goes from 0.2732 to 0.0011.
This variable was not a good predictor in the first model and its coefficient had a low value (0. 48317), even lower than in the control group in the same survey (0.67052). On the opposite, in the second survey its coefficient almost doubles to 1.34776, outclassing the coefficient of the same variable in the control group (0.89785), showing how the treatment group relies on data much more than the control group does.
On the other hand, the coefficient of the independent variable EXEC diminishes in both groups after six months. However, in the treatment group the decrease is less noticeable, thus suggesting that this group may still slightly rely on the executive's opinion during decision-making.
Nonetheless, the high p-value of these coefficients prevents them from providing definitive answers.

5.2.3 Scenario 2

	EXT = 0		EXT = 1	
	Control	Treatment	Control	Treatment
INT = 0	-33.111	2.500	7.417	12.068
INT = 1	-13.150	-8.857	-18.617	20.400

Table 8: Scenario 2: Differences in Means (After - Before)

By analysing the differences in the means between the two surveys, a divide between the two groups' behaviour is immediately noticeable, as shown by table 8.
While the control group increases its preference for external innovations in three out of four versions of the scenario, the treatment group does the opposite. Surprisingly, the latter increases its preference for internal innovations, more so in the versions where the external opportunity is stronger (EXT = 1), although this may merely be an adjustment towards the actual group tendency. As a matter of fact, in the version where EXT and INT are both equal to 1, there is only one respondent to the first survey for the treatment group, so its opinion may not be representative of the whole population.
Moreover, it must also be noted that the treatment group, despite increasing its preference for internal innovations, is still more keen on using external innovations than the control group, as its mean values are generally lower than the treatment group's.
This may indicate an innate preference for external innovations on the treatment group's side, which the control group does not possess with the same strength.

Considering only the control group in the regression analysis, it can be seen in table 9 how in the answers to the second survey the predictive power of the model plummets as its p-value goes from 0.0008 to 0.4117. The same happens to the parameters' significance, so that they are not valid predictors in the second model. Moreover, both parameters EXT and INT become weaker.
These results indicate that the independent variables lose their explanatory power in the passage from the first to the second survey. The changes in the dependent variable, initially easily explainable by the variations in EXT and INT, are now less predictable.
The findings also highlight the fact that the control group in the second survey becomes less reliant on the characteristics of the different innovations when taking a choice.

Considering only the treatment group, the changes from the first to the second survey are even more evident in table 9, in particular because of the low significance of the second model. As

DV : Preference for in-house rather than external innovation				
	Control Group		Treatment Group	
	Before	After	Before	After
Intercept	0.2547	-0.0110	-0.1340	0.1661
	(0.2334)	(0.4407)	(0.3250)	(0.4082)
EXT	-0.7380***	-0.2669	-0.6607	-0.2774
	(0.2551)	(0.4330)	(0.4110)	(0.4593)
INT	0.6692**	0.4056	0.4015	-0.0943
	(0.2545)	(0.4291)	(0.4473)	(0.4593)
Adj. R^2	0.2896	-0.0062	0.1205	-0.0743
F-Value	0.0008	0.4117	0.0820	0.8166
N	39	26	27	24

*p < 0.1, **p < 0.05, ***p < 0.01 ****p < 0.001

Table 9: Scenario 2 Regression Model: NIH Syndrome

when considering other subsamples, the explanatory power of the model decreases, as shown by the lower Adj. R^2 and p-value in table 9.
Another interesting discovery is the fact that the coefficients of the treatment group are in general weaker than the coefficients of the control group, and have less chance to be significant for the determination of the dependent variable.
This means that the treatment group is less reliant than the control group on the characteristics of the different types of innovations during its decision making process. Moreover, as in the control group, this reliance further diminishes after the treatment has been administered.

5.2.4 Scenario 3

	HUM = 0		HUM = 1	
	Control	Treatment	Control	Treatment
MACH = 0	-5.357	-4.500	-7.857	0.500
MACH = 1	-0.181	17.500	6.083	5.000

Table 10: Scenario 3: Differences in Means (After - Before)

The changes in the means of control and treatment groups are quite similar in Scenario 3 except for a couple cases, as seen in table 10. The first is the case where HUM = 0 and MACH = 1, that is to say the version where the human option is not appropriate for the task at hand while the machine option is the most suitable. Six months after the first survey, while the control group strengthens its leaning towards the *do not delegate* option, the treatment group decisively increases its preference for delegating the task to a software, the most appropriate option for this version of the scenario.

Moreover, the control group improves its capacity to identify the most appropriate option for the case where MACH = 0 and HUM = 1, that is to say not delegating the task. This change, however, is smaller than the one the treatment group undergoes.

The results do however underline the fact that both groups after six months improve their ability to identify the most appropriate option for a given vignette, in particular the treatment group.

The means reveal another puzzling fact: the treatment group, compared to the control group, is more likely not to delegate the task when the human option is weaker, and is more likely to delegate when the human option is stronger.

Moreover, as table 14 shows, both groups are decisively more likely to delegate the task when the human is the most appropriate agent for its execution. This seems the opposite of what could be expected to normally happen, although it pertains more to the innate characteristics of the respondents than to the changes brought by the administration of the treatment.

DV : Preference for delegating the task to a machine				
	Control Group		Treatment Group	
	Before	After	Before	After
Intercept	0.4144*	0.2044	0.0816	0.1035
	(0.2162)	(0.3375)	(0.3552)	(0.4161)
MACH	-0.0994	-0.0745	0.4157	0.2964
	(0.3116)	(0.4006)	(0.4170)	(0.4285)
HUM	-0.7148**	-0.3900	-0.7924*	-0.4272
	(0.2813)	(0.3994)	(0.4170)	(0.4285)
Adj. R^2	0.1050	-0.0422	0.0793	-0.0216
F-Value	0.0513	0.61603	0.1420	0.4817
N	39	26	27	24

*p < 0.1, **p < 0.05, ***p < 0.01 ****p < 0.001

Table 11: Scenario 3 Regression Model: Delegating to Machines

Considering only the control group in the regression analysis, the first model in table 11 looks quite reliable, as its p-value hovers at 0.0513. However, only the HUM parameter has high validity (p-value: 0.0155), while the MACH parameter is an inadequate explanator, with a p-value equal to 0.7515.
The MACH coefficient, initially equal to –0.0994, does not undergo any significant changes in the second model, where it is equal to –0.0745, showing how its influence on the dependent variable is null. On the other hand, the HUM parameter becomes weaker, going from a value of –0.7148 to a value of –0.3900.
This shows that the treatment group does not really take into account the characteristics of the software during its decision-making process, but on the contrary carefully analyses the capabilities of the human option.
Moreover, both parameters lose their significance in the second survey, showing how the second model is inadequate to explain the variance of the dependent variable.

Considering only the treatment group in table 11, some peculiarities can be observed when compared to the control group, both in the first and second survey.
First of all, the models built for the treatment group exhibit a low validity in both surveys, not only in the second one like the control group. Thus, it seems likely that the two variables HUM and MACH are not well suited to explain the choices of the treatment group.
However, in the treatment group the values and the validity of the MACH coefficient is always higher than in the control group, thus suggesting that the former tends to take into account the characteristics of the software when making a decision, contrary to what the control group does.
On the other hand, the HUM parameter behaves like in the control group, exhibiting a negative coefficient in both surveys, which weakens, along with the variable's validity, in the second survey.
In general, the treatment group's parameters in table 11 follow the same trend as the control group's after having received the treatment, that is to say they weaken and lose significance. It would seem that both groups' attention to software and human features during the decision-making process has diminished after six months. On the other hand, this decrease could be due to the lower sample size of the second survey, that would not permit the construction of a model that can successfully capture the respondents' tendencies.
In conclusion, in the second survey the model seems to be a better predictor of the dependent variable for the treatment group, although in both groups its significance is low. On the opposite, in the first survey the model is a better predictor for the tendencies of the control group.

6 Discussion

6.1 Correlations

The first interesting finding of this research emerges from the analysis of correlations between dependent variables of different scenarios.

On the one hand, in both surveys a strong correlation can be observed in the control group between the responses to scenario 1, that tests the tendency of relying on data rather than on an opinion, and scenario 3, that tests the tendency of relying on machines rather than on humans for complex tasks.

On the other hand, this correlation is not initially present in the treatment group, but emerges after the treatment has been administered, and in the end it is even stronger than in the control group.

A similar result can be observed in the correlation between the dependent variable of scenario 2 and the dependent variable of scenario 3. The former measures the tendency to rely on internal rather than external innovations, while the latter measures the tendency of relying on a machine rather than on oneself to execute a complex task.

While said correlation is non-existent in the control group in both surveys, the same is not true for the treatment group. In the latter, after the treatment has been administered, a strong and significant inverse correlation appears.

The two results seem to show that only among the students who received the treatment, that is to say attended courses on Big Data and AI, decision-making processes along the variables considered in this study have turned from not being aligned to being aligned, as the respondents have developed strong correlations between the answers to different scenarios.

If we consider the answer to scenario 3 as an indicator of the trust that the students have in machines, it seems that the more the respondents trust machines, the more they also tend to suggest launching a new product, and the more they tend to use external rather than internal innovations.

Moreover, answers to scenario 1 and 2 are never significantly correlated, suggesting two implications. First, it is likely that some feature embodied by the answer to scenario 3 is responsible for the alignment of the decision-making processes. Second, there is no apparent direct link between a preference for basing decisions on data and incorporating external innovation.

The factor on which this mindset alignment is based may be the trust in software, as contemplated before, or the ability to critically compare the capabilities of machines with those

of humans. The latter case would mean that respondents with little knowledge of Big Data and AI do not possess a unique framework for reasoning about the three different issues here investigated, that are instead approached with separate and compartmentalized decision-making processes. On the contrary, students who received the treatment would be able to analyse the different issues with the same mental framework, thus exhibiting correlations between their answers.
It can therefore be concluded that the treatment has had a distinct effect on how the way of thinking of a person is organized, as treated respondents consider the issues presented to them with the same mindset.
This could be the result of a new framework that was taught to the students during the courses they took, or that was absorbed as a consequence of the exposure to these particular topics.
As a result, the treated respondents now approach the different tendencies tested here with a common mindset, giving answers based on common beliefs and ideas as a guide for their behaviour.
This serendipitous finding shows how the courses have indeed reached some of their objectives, such as teaching students a framework to critically analyse questions and issues related to the use of data (CBS, 2017a, 2017b; Bocconi, 2017).

6.2 Scenarios

6.2.1 Scenario 1

The results from Scenario 1 seem to indicate that the treatment has had an effect.
After having received the treatment, the DATA variable shows a decisively bigger change in the treatment group than in the control group.
In the former, while in the first survey DATA is not a good predictor of the answer and its influence is moderate, it becomes an excellent predictor in the second model, and its coefficient increases almost threefold, to 1.34776. Thus, the treatment group in the second survey is much more reliant on data than the control group when taking decisions, since a change from unfavourable to favourable data results in a much higher likelihood to suggest the launch of the product.
Furthermore, looking at the difference in the changes of this parameter the contrast between control and treatment group becomes even more apparent. The change in the DATA coefficient from survey 1 to survey 2 in the case of the control group is 0.227, while in the

treatment group is more than three times that (0.865).
These results seem to suggest that the treatment has had the effect of increasing the reliance on data during the decision-making process of the respondents.
The fact that attending a course on Big Data or AI increases one's reliance on data does not come as a surprise, as it confirms previous findings (Dunlap & Piro, 2016; Kippers et al., 2018), but it is nonetheless important in proving that this kind of university courses has the effect of teaching students to rely more on data during their decision-making process.

On the other hand, in regard to the tendency to rely on HiPPO and confirmatory bias, the results appear mixed. In the first survey both groups seem to be influenced by a superior's opinion, although without much significance, particularly when it comes to the treatment group.
As a matter of fact, changing the executive's opinion in regard to the product's success seems to have no effect on the answer of the respondents, possibly because this opinion is never taken into consideration.
The low validity of EXEC in the first survey could mean that both groups, and the treatment group in particular, have a pre-existent tendency to not take a HiPPO into consideration when taking a decision based on data, and at the same time exhibit a low confirmatory bias.
This tendency is even more apparent in the treatment group during the second survey, as the validity of EXEC decreases, showing that it is even more likely for the dependent variable not to be affected by the executive's opinion.
This behaviour can also be observed in the control group, where the EXEC variable's significance drops after six months.
The fact that both groups exhibit this tendency in the second survey may bring to two possible conclusions.
First, something may have occurred in the six months that separate the two surveys other than the treatment, and has resulted in a general reduced trust in HiPPO across the whole sample.
Second, the second survey did not accurately capture the actual tendencies of the sample in regard to the EXEC variable.
The second option is certainly possible, but it may be contradicted by the fact that the models for both the treatment and control group appear quite solid, with a p-value of 0.0678 and 0.0041 respectively. Still, it must be noted that this may be due mostly to the high reliability of the DATA parameter for explaining the dependent variable. Thus, it is also possible that the EXEC parameter is not adequately captured by the model. This may be due to the sample size or to the vignette formulation.

The analysis of the means reveals a preference of the respondents for suggesting the launch of the product, as the mean value of the responses is always above 50, the value that expresses neutrality.
This tendency is however reduced in the second survey, as the mean values of the responses decrease. This change is more evident in the treatment rather than in the control group: while in the latter the overall mean decreases only by 1 point, in the former it does so by 6 points.
This finding is further explained by the analysis of the means for each different version of Scenario 1. As table 12 shows, in the case where the executive is confident about the launch of the product but the data do not suggest success, after six months respondents have increased their ability to select the most appropriate course of action. As a matter of fact, they are much more likely than before to advise against a product launch. This change is more pronounced in the case of the treatment group, as the mean value of the responses for this scenario version decreases by 30 points, as opposed to a 15 points decrease in the control group.
A similar result transpires from the answers to the version where both the executive and the data suggest success. As in the other version, the treatment group increases significantly more than the control group its ability to suggest the most appropriate course of action.
Thus, a clear result emerges: the treatment group improves more than the control group its ability to correctly assess the best choice for a given vignette, as its answers become more dependent on the actual data of the scenario rather than on the opinion of the executive.

The results of both regression and mean analysis prove that the treatment affects the students in a positive way. Those who receive it become more reliant on data for their decisions, and tend to select the best option for a given vignette.
This result confirms the finding of previous studies on data literacy (Dunlap & Piro, 2016; Kippers et al., 2018), showing how exposure to in-depth knowledge about data can increase a person's ability to use data effectively.

In conclusion, ***Sub-Question 1* can be answered positively**, as students rely more on data rather than opinions after the treatment has been administered.

6.2.2 Scenario 2

The results of Scenario 2 provide an interesting outlook of respondents' tendencies in the first survey and the way they change in the second survey, although the findings are overall mixed and do not permit a definitive answer to *Sub-Question 2* .

In general, the regression model built from the results of the first survey shows how a better external opportunity makes respondents more likely to choose to incorporate an external innovation within the company. Similarly, a better internal opportunity makes them more likely to choose to develop the innovation in-house. These results are indeed not surprising, as they only show that respondents do take into account the characteristics of both types of innovation during their decision-making process.

Another common characteristic of the two groups is that the influence of INT is slightly weaker than the influence of EXT on the decision of the respondent, suggesting that respondents may care less about the characteristics of the internal innovation, possibly because they are more critical of the external innovation and scrutinize it more thoroughly.

This result also seems to disprove the fact that there is a preexisting bias towards innovations developed in-house, much documented in previous literature (Katz & Allen, 1982; Wastyn & Hussinger, 2011), as if that were the case the INT parameter would have been expected to be stronger.

The existence of a NIH Syndrome in the respondents is further disproved by the results of the mean analysis. As table 13 shows, the respondents are generally more likely to choose the external innovation over the innovation developed in-house. This tendency is especially noticeable in the treatment group, which in almost all versions of the scenario shows a greater preference for external innovations compared to the control group.

Some research does indeed support the absence of a bias for internal innovations, as it has been shown that some companies may have a positive bias towards external knowledge (Piller & Antons, 2015).

Nonetheless, the respondents to these surveys are students, not company managers or even company employees. Without having worked in a company for some time, they have most likely not yet experienced a sense of identification with it, and they almost certainly do not have strong feelings for the company described in the survey.

Therefore, it may be that NIH syndrome does not manifest itself in the respondents simply because they lack identification with the company or a group within it, a critical trait for the development of NIH behavior (Wastyn & Hussinger, 2011).

A more puzzling result emerges from the changes between the first and second survey. Across

all respondents the coefficients and the validity of the dependent variables plummet, highlighting a general decrease in the relevance of the innovations' features for the choice of which innovation to use.
This trend is more pronounced in the treatment group, that after the treatment's administration appears to lose all consideration for the characteristics of the internal innovation during its decision-making process.
Therefore, it may seem that after six months the dependent variables have lost their influence on the decision-making process of the respondents. However, this drop in significance and influence can also be explained by the smaller sample size of the second survey, that may not allow to accurately capture the tendencies of the respondents about NIH Syndrome.

In general, respondents seem to become more moderate in their preferences, leaning more towards neutral values in their answers after six months have passed. As this phenomenon is evident in both groups, it does not seem likely for the treatment to have had an effect on the students.
In the end, the answers to this scenario ultimately lead to two relevant results. First, the students in the treatment group show an innate preference for external innovations, not present in the control group. Second, six months after the first survey, respondents are less certain about which option to choose, and the features of the innovations are less relevant to their decision.
There is however no indication of an increase in the treatment group's preference for external innovation.

In conclusion, ***Sub-Question 2*** **cannot be answered positively**.

6.2.3 Scenario 3

The results of scenario 3 do not provide significant findings in regard to the efficacy of the treatment for the purpose of investigating *Sub-Question 3*, but they manage to uncover other intriguing changes in the preferences of the respondents.
In general, there are only a few differences between control and treatment group in how the reliance on human and machine features changes between the two surveys.
As already described in Scenario 2, the strength of the coefficients wanes after six months, highlighting a decrease in the reliance that the respondents have on both human and software characteristics during their decision-making process.

In particular the HUM parameter, a good predictor of the dependent variable for both groups in the first survey, loses much of its significance and strength in the models for the second survey. However, the fact that it is a valid and strong predictor for the dependent variable in the first survey shows that respondents do take into account the characteristics of a human decision-maker when considering whether to delegate a task, at least initially.

On the other hand, it is impossible to ascertain that the MACH parameter ever has a significant influence the dependent variable, in particular for the control group. The software features do nonetheless seem to play a role in the decision-making process of the treatment group, as the value of its coefficient is above or equal to 0.3 in both surveys. On the opposite, the respondents in the control group never consider the characteristics of a machine during their choice of whether to delegate a task to it.

These results seem to indicate that the choice to delegate is still heavily based on the evaluation of the skills of a person and less on the skills of a machine. Previous research has found that change in delegation behaviour is caused by ICT and the tools to share information, rather than the tools to analyse it (Bloom, Garicano, et al., 2014; Colombo & Delmastro, 2004). Although these concepts are inherently linked, the technologies that make up the treatment lean more on the latter, and the results do not show any increase in delegation. This further suggests that the delegation phenomenon previously observed (Hitt & Brynjolfsson, 1997) concerns human, not software, capabilities. If the development of communication, rather than decision-making tools, is the factor that contributes to an increase in delegation, that simply means that humans are still the agents taking decisions. As a matter of fact, delegation may increase merely because it becomes easier for a person to access knowledge from different sources and be empowered to take decisions, showing how human skills are still the main factor considered during a delegation decision.

Therefore, it seems that, as shown by previous research (Hitt & Brynjolfsson, 1997), recent developments in technology have not eroded the traditional human tendency to look first and foremost for the right person for the job.

The treatment group, however, seems to possess a pre-existing tendency to more easily take software features into consideration when choosing whether to delegate a task to a machine. This may be due to the fact that students who have chosen to take part in courses which main topics concern the capabilities of computers also probably have an interest in computers and software, that can make them more receptive to the features of a software. Thus, they may be able to more accurately discern the advantages and disadvantages of different software solutions.

This possibility, however, cannot be confirmed as a consequence of the low significance of the model in the second survey.

Nonetheless, the self-selected treatment group seems to possess an overall greater awareness of the importance that should be given to both the characteristics of the human and the machine when taking a decision. On the opposite, the control group seems to only consider the human factor, while regarding the differences in software features as unimportant.

As with the coefficients of the regression model, the analysis of the means does not show many differences between control and treatment groups in the changes between first and second survey.
Two are the only notable differences in the way the means change after six months have passed. First, the treatment group greatly increases its preference for delegating when the machine possesses more suitable features than the human for the task at hand. Second, the control group increases its preference for not delegating when the software is less appropriate than the human for executing the task. In the other cases, the changes in the means are small and similar across groups.
Likewise, when it comes to group preferences, there are no vast differences between the attitudes of the two groups. In general, students lean decisively towards the option *delegate to a software*, showing perhaps a certain trust in the ability of machines to complete complex tasks.
Furthermore, both groups do not seem to be able to always correctly discern which option is most appropriate for a given version of the scenario, as they prefer delegating the task even if not doing so is the most appropriate choice. This reflects the findings of the regression analysis, and further shows how the characteristics of the different options become less important for the respondents' final choice.
Both groups, however, do improve their ability to take the most appropriate choice in the second survey, with no increase in their preferences for delegating to machines. Since there are no significant differences between the groups, this change cannot be attributed to the administration of the treatment.

In conclusion, ***Sub-Question 3*** **cannot be answered positively**, but other potentially interesting findings emerge from the results.
As this hypothesis was the one least based on the existing theory, this result may not be surprising.

However, the fact that the administration of the treatment has no significant effects on the respondents should not be dismissed as unimportant. If combined with the findings from the other scenarios and the correlation analysis, a noteworthy pattern can be observed in the treated students.

6.3 Connecting the Dots

The final results of the three scenarios seem to suggest that the treatment does indeed have no significant effects on the tendencies considered, besides increasing data literacy.
This should not, however, be taken as a sign that the findings are inconclusive. On the contrary, the inability to answer positively to sub-questions 2 and 3 clears the way for important considerations on the effects of the treatment.

On the one hand, the results of correlation analysis and the findings of Scenario 1 show that the treatment has had an effect. Using and learning how to use Big Data and advanced analytical tools not only changes a student's mindset and her approach to issues involving machines, but also increases her reliance on data during decision-making.
On the other hand, the same treatment seems to have no effects on the respondents' tendencies concerning machines and external innovation. As the effect on the respondents' mindset and data attitudes is quite evident, this may come as a surprise.
This result, however, shows that students who get educated in the use of Big Data, AI, and new technologies, manage to acquire data literacy without losing respect for and confidence in human labour. This change is most likely a positive one. As automation steadily creeps into the workplace, being able to recognize and harness human contribution for problem solving, as the respondents do, will be an essential skill for managers. As a matter of fact, the automation of many tasks now performed by humans will not be a sudden event, but a slow process that will require the contribution of humans at every stage, despite what recent articles may suggest (Arntz, Gregory, & Zierahn, 2016). Therefore, quickly dismissing the importance of human contribution and skills would be a severe miscalculation.
It is indeed easy to think that a machine will soon be able to execute most tasks that are now the domain of humans (Barrat, 2015), despite the fact that even the newest technologies have several limitations on the tasks they can perform (Kelly, 2017). It would therefore be troubling if managers and students exposed to AI and Big Data were to favour machines without reserve during their decision-making process, as they could easily incur in automation bias (Skitka, Mosier, & Burdick, 2000; Wickens, Clegg, Vieane, & Sebok, 2015), and would not be using the best resource for completing a given task.
In fact, human labour seems to often complement automation, and the current and near-future abilities of machines should not be overestimated, a mistake that could result in a

dangerous misallocation of resources (Autor, 2015).
Thankfully, the treatment does not seem to reflect this possibility, and has instead what one may consider mainly positive effects on students, as it broadens their data analysis capabilities without making them neglect the power of human contribution.
Therefore, it seems that the courses considered are effectively teaching students the right competences needed to excel in the workplace and in their career. As they will likely be the managers of the future, the thought that they will be able to value human creativity and skills is certainly welcomed news, and this finding can contribute to the ongoing discourse about automation and the future of human labour (Brynjolfsson & McElheran, 2016; Arntz et al., 2016).

7 Conclusions

The aim of this research was to investigate whether exposure to technologies powered by Artificial Intelligence and Big Data can affect a person's decision-making process.
To do so, business students taking their master's degree were surveyed about their decision-making tendencies before and after some of them took at least one elective course regarding the aforementioned technologies, although most of them enrolled in three or four courses on the topic. The changes in their attitudes were then analysed to see whether the in-depth exposure to these topics had affected their decision-making process, with an experiment design taking inspiration from difference-in-differences models.
The students who took the courses were considered the treatment group, and the ones who did not, the control group.
The main research hypothesis of this study was that the treatment would have additional effects on the sample's decision-making tendencies other than increasing data literacy. Subsequently, three decision-making tendencies were tested. First, data literacy, treated as reliance on HIPPOs and confirmation bias rather than data. Secondly, the presence and strength of a Not-Invented-Here Syndrome. Lastly, the tendency to delegate tasks to machines.
The treatment has been found to have an effect on the first tendency, as students who received it were drastically more reliant on data than before, when compared to the control group. On the other hand, no significant results could be found in regard to the ability of the treatment to have an effect on the other two tendencies.
Nonetheless, these findings can be highly informative about the tendencies of the surveyed sample.
For instance, the results of the scenario dealing with NIH syndrome showed that students do not inherently exhibit this kind of bias, a not so self-evident finding. Furthermore, the results of the scenario testing the tendency to delegate to a machine suggested that students who take courses on how to use data do not become blinded by their perks, and do instead recognise the important contribution that human work can make.
There were, moreover, unexpected findings during the analysis of the correlations between the answers to different scenarios. The tendencies of the treatment group, previously independent from each other, became correlated after the treatment had been administered.
This suggests that the courses may have had the effect to stabilize and standardize the students' ways of thinking about these topics.
If these results are considered together, the treatment is found to increase respondents' reliance on data, while at the same time preventing the development of an automation bias or

a decrease in the quality of the respondents' assessment of human work.
The implications of these findings can be divided in two broad categories, depending on which entity, namely companies and academic institutions, is most affected by them.

7.1 Implications for Academia

The results of this research show that the courses taken by the students reach their aim of having an effect on students' mindsets and making them capable of critically analysing data. After taking the courses, students show a common framework for thinking about decision-making situations associated with concepts such as data, open innovation and delegation to machines, in addition to being more aware of the characteristics of the data and software they are working with.
Additionally, the fact that the treatment increases students' data literacy, while at the same time making them skeptical of blindingly trusting data, can be seen as a sign that the courses have at least in part reached their learning objectives (Bocconi, 2017; CBS, 2017a, 2017b).
This study thus provides an indicator for understanding the value that these courses can offer, so that students can take an informed choice when deciding whether to enrol in them. Moreover, it also offers academia a way to assess the impact of higher education level courses on students who attend them.
The fact that students develop a particular mindset after attending the courses considered may also open a wider discussion about the merits of teaching such a mindset. It also highlights the importance of carefully considering the concepts and frameworks that are taught, since these will undoubtedly have an influence on the way students think and take decisions.

7.2 Implications for Business

The findings can also offer valuable indications for businesses.
First of all, the students in the treatment group seem to be more able to critically discern different software and data characteristics while taking a decision when compared to the control group, in particular after the treatment has been administered.
Companies should then be aware that students taking courses on AI and Big Data are more capable than others to understand and critically interpret data analytics. As these students will likely be the managers of tomorrow, and managers who understand data analytics have

a greater professional advantage over those who do not (Migliore & Hubbard, 2016), companies can benefit from this advantage, also thanks to the fact that these people can effectively harness data-driven decision-making to boost performance (Brynjolfsson & McElheran, 2015, 2016).

Secondly, this research provides further confirmation that business students, future employees, do not have a pre-existing NIH Syndrome. On the contrary, they seem to welcome the incorporation of external innovations into internal product development.

This is important information for companies, that can thus act to assure that the syndrome is never developed. This can be done by devising incentives for encouraging the adoption of external innovation, and by promoting a culture that does not assume in-house superiority, like some companies are already doing (Piller & Antons, 2015).

Lastly, some of the findings seem to suggest that the students of the institutions considered are not afraid of delegating complex tasks to software but at the same time do not forget the importance of human work. This should be welcomed as good news by business, as it means that there exists a workforce capable of working in harmony with machines without discriminating against humans. Since many of the new jobs created by automation will require these types of skills (The Economist, 2017a; Brynjolfsson & Mcafee, 2014), companies should be aware of the benefits that such a workforce can entail.

7.3 Limitations

This study presents various limitations.

The first issue pertains to the sample selection. Although I believe business students at the end of their academic experience are a close enough representation of managers, as it is often the case in research (Hutchinson et al., 2010; Schulz-Hardt et al., 2000; Aiman-Smith et al., 2002), they could differ in some characteristics that may influence their responses. However, the use of students as a sample allowed this research to produce findings about the effectiveness and the effects of some of the courses taught at Bocconi and CBS.

A second limitation is the sample size, that could have resulted in the non-significance of some of the regression models and may have limited the validity of the findings. This limitation was a direct result of the treatment design, as there is a limited number of students taking the kinds of courses appropriate to be considered a valid treatment.

Another related unwanted influence could have taken place because of the different formulations of the scenarios in the two surveys, which may differ in some key points so that they do not present a similar enough situation across both surveys. That is to say, the same scenario presented in two different surveys may have not been similar enough to affect the respondents' answers in the same way, thus eliciting and measuring slightly dissimilar tendencies. These limitations, however, do not significantly reduce the validity of the results obtained by the experiment.

7.4 Directions of Future Research

The findings of the present research, together with its limitations, can certainly offer some indications on other interesting lines of research that would be worth exploring by future studies.

The most obvious might be the testing of the same hypotheses exposed in this thesis with a larger sample, or with a sample whose respondents are managers instead of students. In case the experiments were to be repeated with students, it would certainly be beneficial to include students from more courses and more universities, such as all those which share a Double Degree agreement with EMIT.

Another important development of the work presented here would be finding a solid indicator of trust in machines and explore the potential correlations and causal relationships that it may have with the decision-making tendencies considered in this study.

Moreover, finding out how trust in machines and software is born and which factors influence its development would surely be a related interesting topic of research.

Furthermore, there may be room for the analysis of NIH Syndrome-related tendencies that can affect its emergence, such as the way communication is handled (Katz & Allen, 1982) or the way knowledge is organized (Husted & Michailova, 2002). Devising a vignette study where these elements are the independent variables and students are the respondents could shed light on whether no NIH Syndrome is actually present or if this research did not take into account the factors that promote its appearance.

Lastly, it would be intriguing to investigate more extensively the attitudes of future and current managers about human and machine workers. While this thesis can provide a starting point for a deeper analysis of the relationship between human and machine contribution to problem-solving, it is certainly not a comprehensive study of companies' attitudes about the issue.

This research has provided a first outlook at the way the latest analytics technologies shape some particular aspects of the human mind related to the decision-making process.

The limitations of this study will hopefully be a starting point for new research, and its findings hints for a wider discussion about the role that students, managers, data and universities can play in a world increasingly centered around smarter machines.

8 References

Agrawal, A. K., Cockburn, I. M., Audretsch, D., Jarmin, R., Glaeser, E., Rosenthal, S., … Thompson, P. (2010). Not Invented here? Innovation in Company Towns. *Journal of Urban Economics*, *67*(1), 78–89.

Aguinis, H. & Bradley, K. J. (2014). Best Practice Recommendations for Designing and Implementing Experimental Vignette Methodology Studies. *Organizational Research Methods*, *17*(4), 351–371. Retrieved from http://journals.sagepub.com/doi/10.1177/1094428114547952

Aiman-Smith, L., Scullen, S. E., & Barr, S. H. (2002). Conducting Studies of Decision Making in Organizational Contexts: A Tutorial for Policy-Capturing and Other Regression-Based Techniques. *Organizational Research Methods*, *5*(4), 388–414. Retrieved from http://journals.sagepub.com/doi/10.1177/109442802237117

Arntz, M., Gregory, T., & Zierahn, U. (2016). *The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis*, OECD. Paris. doi:10.1787/5jlz9h56dvq7-en

Atzmüller, C. & Steiner, P. M. (2010). Experimental Vignette Studies in Survey Research. *Methodology*, *6*(3), 128–138. doi:10.1027/1614-2241/a000014

Autor, D. H. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives*. doi:10.1257/jep.29.3.3

Baesens, B., Bapna, R., Marsden, J. R., Vanthienen, J., & Zhao, J. L. (2016). Transformational Issues of Big Data and Analytics in Networked Business. *Mis Quarterly*, *40*(4), 807–818. doi:10.5121/ijgca.2012.3203

Barrat, J. (2015). Our Final Invention : Artificial Intelligence and the End of the Human Era. doi:10.1017/CBO9781107415324.004

Beer, D. (2016). How should we do the history of Big Data? *Big Data & Society*, *3*(1), 205395171664613. doi:10.1177/2053951716646135

Benson, B. (2016). Cognitive Bias Cheat Sheet. Retrieved from https://betterhumans.coach.me/cognitive-bias-cheat-sheet-55a472476b18

Bertrand, M., Esther, D., & Sendhil, M. (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*, *119*(1), 249–275.

Bishop, P. A. & Herron, R. L. (2015). Use and Misuse of the Likert Item Responses and Other Ordinal Measures. *International journal of exercise science*, *8*(3), 297–302. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/27182418%5Cnhttp://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4833473

Bloom, N., Brynjolfsson, E., Foster, L., Jarmin, R., Patnaik, M., Saporta-Eksten, I., & Van Reenen, J. (2014). CEP Discussion Paper No 1258 - IT and Management in America. *Centre for Economic Performance*, (1258).

Bloom, N., Garicano, L., Sadun, R., & Van Reenen, J. (2014). The Distinct Effects of Information Technology and Communication Technology on Firm Organization. *Management Science*, *60*(12), 2869–2885.

Bocconi. (2017). 20564 - BIG DATA FOR BUSINESS DECISIONS. Retrieved from http://didattica.unibocconi.it/ts/tsn_anteprima2006.php?cod_ins=20564&anno=2018&IdPag=6067

Boncea, R., Ionut, P., Smada, D.-M., & Zamfiroiu, A. (2017). A Maturity Analysis of Big Data Technologies. *Informatica Economica*, *21*(1/2017), 60–71. doi:10.12948/issn14531305/21.1.2017.05

Boudreau, D. (2017). Five Ways in Which Artificial Intelligence Changes the Face of Web Accessibility. Retrieved from https://www.24a11y.com/2017/five-ways-in-which-ai-changes-the-face-of-web-accessibility/

Boyd, D. & Crawford, K. (2012). Critical Questions for Big Data: Provocations for a Cultural, Technological and Scholarly Phenomenon. *Information, Communication & Society*, *15*(5), 662–679. doi:10.1126/science.1243089

Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence. *The Quarterly Journal of Economics*, (February), 339–377. doi:10.1162/003355302753399526

Brynjolfsson, E. & Hitt, L. M. [Lorin M]. (1996). Paradox Lost? Firm-level Evidence on the Returns to Information Systems Spending. *Management Science*, *42*(4), 541–558.

Brynjolfsson, E. & Hitt, L. M. [Lorin M.]. (1998). Beyond the Productivity Paradox: Computers are the Catalyst for Bigger Changes. *Communications of the ACM*, *41*(8), 49–55. doi:10.1145/280324.280332

Brynjolfsson, E., Hitt, L. M., & Kim, H. H. (2011). Strength in Numbers: How Does Data-Driven Decision Making Affect Firm Performance? *SSRN Electronic Journal*. doi:10.2139/ssrn.1819486

Brynjolfsson, E. & McAfee, A. (2013). The Great Decoupling. *New Perspectives Quarterly*, *30*(1), 61–63. doi:10.1111/npqu.11362

Brynjolfsson, E. & McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (First). New York: W. W. Norton & Company.

Brynjolfsson, E. & Mcafee, A. (2014). The digitization of just about everything. *The second machine age: work, progress, and prosperity in a time of brilliant technologies*, 57–70.

Brynjolfsson, E. & McElheran, K. (2015). Data in Action : Data-Driven Decision Making in US Manufacturing, 1–55.

Brynjolfsson, E. & McElheran, K. (2016). The Rapid Adoption of Data-Driven Decision-Making. *American Economic Review*, *106*(5), 133–139. doi:10.1257/aer.p20161016

CBS. (2017a). KAN-CCMVV2556U Big Data Analytics. Retrieved from http://kursuskatalog.cbs.dk/2017-2018/KAN-CCMVV2556U.aspx

CBS. (2017b). KAN-CCMVV4042U Datafication – foundations, transformations and challenges. Retrieved from http://kursuskatalog.cbs.dk/2017-2018/KAN-CCMVV4042U.aspx?lang=en-GB#

Chen, H. & Storey, V. C. (2012). Business Intelligence and Analytics : From Big Data To Big Impact. *Mis Quarterly*, *36*(4), 1165–1188. doi:10.1145/2463676.2463712

Clagett, R. P. (1967). *Receptivity of Innovation - Overcoming N.I.H.* (Doctoral dissertation, MIT).

Colombo, M. G. & Delmastro, M. (2004). Delegation of Authority in Business Organizations : An Empirical Test. *The Journal of Industrial Economics*, *52*(1), 53–80. Retrieved from http://www.jstor.org/stable/3569762

Côrte-Real, N., Oliveira, T., & Ruivo, P. (2017). Assessing business value of Big Data Analytics in European firms. *Journal of Business Research*, *70*, 379–390. doi:10.1016/j.jbusres.2016.08.011

Davenport, T. H., Barth, P., & Bean, R. (2012). How big data is different. *MIT Sloan Management Review*, *54*(1), 43.

David, P. A. (2016). American Economic Association The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *The American Economic Review*, *80*(2), 355–361.

Demchenko, Y., Grosso, P., & Membrey, P. (2013). Addressing Big Data Issues in Scientific Data Infrastructure. In *Collaboration technologies and systems (cts), 2013 international conference on* (pp. 48–55). IEEE.

Dickson, B. (2017). Artificial Intelligence Creates New Job Opportunities. *PC Magazine*.

Dos Santos, B. L., Peffers, K., & Mauer, D. C. (1993). The Impact of Information Technology Investment Announcements on the Market Value of the Firm. *Information Systems Research*, *4*(1), 1–23. doi:10.1287/isre.4.1.1

Dunlap, K. & Piro, J. S. (2016). Diving into data: Developing the capacity for data literacy in teacher education. *Cogent Education*, *3*(1). doi:10.1080/2331186X.2015.1132526

Ekbia, H., Mattioli, M., Kouper, I., Arave, G., Ghazinejad, A., Bowman, T., . . . Sugimoto, C. R. (2015). Big data, bigger dilemmas: A critical review. *Journal of the Association for Information Science and Technology*, *66*(8), 1523–1545.

Fan, W. & Bifet, A. (2013). Mining Big Data : Current Status , and Forecast to the Future. *ACM SIGKDD Explorations Newsletter*, *14*(2), 1–5. doi:10.1145/2481244.2481246

Finch, J. (1987). The Vignette Technique in Survey Research. *Sociology*, *21*(1), 105–114.

Galeon, D. (2017). The World's First Album Composed and Produced by an AI Has Been Unveiled. Retrieved from https://futurism.com/the-worlds-first-album-composed-and-produced-by-an-ai-has-been-unveiled/

Gopalkrishnan, V., Steier, D., Lewis, H., & Guszcza, J. (2012). Big data, big business: bridging the gap. *BigMine '12: Proceedings of the 1st International Workshop on Big Data, Streams and Heterogeneous Source Mining*, *August*, 7–11. doi:10.1145/2351316.2351318

Günther, W. A., Rezazade Mehrizi, M. H., Huysman, M., & Feldberg, F. (2017). Debating big data: A literature review on realizing value from big data. *The Journal of Strategic Information Systems*, *26*(3), 191–209. doi:10.1016/j.jsis.2017.07.003

Habegger, J. (2005). *Sourcebook of Modern Furniture* (Third Edit). New York, New York: W. W. Norton & Company.

Haselton, M. G., Nettle, D., & Andrews, P. W. (2015). The Evolution of Cognitive Bias. In *The handbook of evolutionary psychology* (pp. 724–746). doi:10.1002/9780470939376.ch25

Hitt, L. M. [Lorin M.] & Brynjolfsson, E. (1997). Information Technology and Internal Firm Organization: An Exploratory Analysis. *Journal of Management Information Systems*, *14*(2), 81–101. doi:10.1080/07421222.1997.11518166

Hoc, J.-M. (2000). From human-machine interaction to human-machine cooperation. *Ergonomics*, *43*(7), 833–843. doi:10.1080/001401300409044

Husted, K. & Michailova, S. (2002). Diagnosing and Fighting Knowledge-Sharing Hostility. *Organizational Dynamics*, *31*(1), 60–73. doi:10.1016/S0090-2616(02)00072-4

Hutchinson, J. W., Alba, J. W., & Eisenstein, E. M. (2010). Heuristics and Biases in Data-Based Decision Making: Effects of Experience, Training, and Graphical Data Displays. *Journal of Marketing Research*, *47*(4), 627–642. doi:10.1509/jmkr.47.4.627

Issa, H., Ting, S., & Vasarhelyi, M. A. (2016). Research Ideas for Artificial Intelligence in Auditing: The Formalization of Audit and Workforce Supplementation. *Journal of Emerging Technologies in Accounting*, *13*(2), 1–20. doi:10.2308/jeta-10511

Jonas, E., Schulz-Hardt, S., Frey, D., & Thelen, N. (2001). Confirmation bias in sequential information search after preliminary decisions: An expansion of dissonance theoretical research on selective exposure to information. *Journal of Personality and Social Psychology*, *80*(4), 557–571. doi:10.1037//0022-3514.80.4.557

Katz, B. R. & Allen, T. J. (1982). Investigating the Not Invented Here (NIH) svndrome : A look at the performance , tenure , and communication patterns of 50 R & D Project Groups. *R&D Management*, *12*(1), 7–20.

Kelly, K. (2017). The Myth of a Superhuman AI. Retrieved from https://www.wired.com/2017/04/the-myth-of-a-superhuman-ai/

Kimble, C. & Milolidakis, G. (2015). Big Data and Business Intelligence: Debunking the Myths. *35*(1), 23–34. doi:10.1002/joe.21642

Kippers, W. B., Poortman, C. L., Schildkamp, K., & Visscher, A. J. (2018). Data literacy: What do educators learn and struggle with during a data use intervention? *Studies in Educational Evaluation*, *56*(July 2017), 21–31. doi:10.1016/j.stueduc.2017.11.001

Kiulian, A. (2017). 6 AI-Powered Tools All Entrepreneurs Need to Run Their Business. Retrieved from https://www.entrepreneur.com/article/298628

Kurzweil, R. (2006). *The Singularity is Near: When Humans Transcend Biology*. Penguin Books.

Lohr, S. (2012). The Age of Big Data. Retrieved from http://rkbusinessschool.com/uploads/2_The_New_York_Times_on_The_Age_of_Big_Data.pdf

Luca, M. (2014). Were OkCupid's and Facebook's experiments unethical? Retrieved from https://hbr.org/2014/07/were-okcupids-and-facebooksexperiments-%0Aunethical

Lucas, L. & Waters, R. [Richard]. (2018). The AI arms race: China and US compete to dominate big data. Retrieved from https://www.ft.com/content/e33a6994-447e-11e8-93cf-67ac3a6482fd

Mahmood, A. M. & Mann, G. J. (2015). Measuring the Organizational Impact of Information Technology Investment: An Exploratory Study. *Journal of Management Information Systems*, *10*(1993), 97–122.

Mandinach, E. B., Honey, M., & Light, D. (2006). A Theoretical Framework for Data-Driven Decision Making. In *Annual meeting of the american educational research association* (pp. 1–18). San Francisco, CA.

Manovich, L. (2001). *The language of new media*. MIT press.

Marr, B. (2017). Instead Of Destroying Jobs Artificial Intelligence (AI) Is Creating New Jobs In 4 Out Of 5 Companies. Retrieved from https://www.forbes.com/sites/bernardmarr/2017/10/12/instead-of-destroying-jobs-artificial-intelligence-ai-is-creating-new-jobs-in-4-out-of-5-companies/#7ee453ea120d

Mayer-Schonberger, V. & Cukier, K. (2014). *Big Data: A Revolution That Will Transform How We Live, Work, and Think* (Reprint Ed). Eamon Dolan/Mariner Books; Reprint edition (March 4, 2014).

McAfee, A. & Brynjolfsson, E. (2012). Big Data: The Management Revolution. *Harvard Business Review*, *90*(10), 60–68. doi:10.1007/s12599-013-0249-5

McAfee, A. & Brynjolfsson, E. (2016). Human Work in the Robotic Future. *Foreign Affairs*, *95*(4), 139–150. doi:10.1525/sp.2007.54.1.23.

McCorduck, P. (2004). *Machines Who Think: A Personal Inquiry Into The History And Prospects Of Artificial Intelligence* (2nd ed.). Ak Peters Series. A.K. Peters. Retrieved from https://books.google.de/books?id=aH9QAAAAMAAJ

McElheran, K. (2014). Delegation in Multi-Establishment Firms: Evidence From IT Purchasing. *Journal of Economics & Management Strategy*, *23*(2), 225–258.

McElheran, K. & Brynjolfsson, E. (2016). The Rise of Data-Driven Decision Making Is Real But Uneven, 2–5. Retrieved from http://mitsloanexperts.mit.edu/the-rise-of-data-driven-decision-making-is-real-but-uneven/

Migliore, L. A. (2017). Demystifying the Big Data Phenomenon for Strategic Leadership. *SAM Advanced Management Journal*, *82*(1), 48.

Migliore, L. A. & Chinta, R. (2016). Mobile technology and the employee-customer-profit chain. *SAM Advanced Management Journal*, *81*(1), 52.

Migliore, L. A. & Hubbard, H. (2016). Fast Learning and Leading in All Things Digital. In *Clute international academic business conference*. Orlando, FL.

Newell, S. & Marabelli, M. (2015). Strategic opportunities (and challenges) of algorithmic decision-making: A call for action on the long-term societal effects of 'datafication'. *Journal of Strategic Information Systems*, *24*(1), 3–14. doi:10.1016/j.jsis.2015.02.001

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, *2*(2), 175–220. Retrieved from http://doi.apa.org/getdoi.cfm?doi=10.1037/1089-2680.2.2.175

Norvig, P. & Russell, S. (2009). *Artificial Intelligence: A Modern Approach* (3rd). Pearson.

Parasuraman, R. & Riley, V. (1997). Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors*, *39*(2), 230–253.

Piller, F. T. & Antons, D. (2015). Opening the Black Box of " Not Invented Here ": Attitudes, Decision Biases, and Behavioral Consequences. *Academy of Management Perspective*, *29*(2), 193–217. doi:10.5465/amp.2013.0091

Pogue, D. (2018). Is Art Created By AI Really Art? Retrieved from www.scientificamerican.com/article/is-art-created-by-ai-really-art/

Provost, F. & Fawcett, T. (2013). Data Science and its Relationship to Big Data and Data-Driven Decision Making. *Big Data*, *1*(1), 51–59. doi:10.1089/big.2013.1508

Qin, J. & D'Ignazio, J. (2010). Lessons learned from a two-year experience in science data literacy education. *International Association of Scientific and Technological University*

Libraries, 31st Annual Conference, 188–204. Retrieved from http://docs.lib.purdue.edu/iatul2010/conf/day2/5%5Cnhttp://onesearch.uoregon.edu/metasearch/record?group=2012-11-04-001379&resultSet=007769&startRecord=1

Rabin, M. & Schrag, J. L. (1999). First Impressions Matter: A Model of Confirmatory Bias. *The Quarterly Journal of Economics*, *1114*(1), 37–82. Retrieved from http://www.jstor.org/stable/2586947

Reed, S., Oord, A. v. d., Kalchbrenner, N., Colmenarejo, S. G., Wang, Z., Belov, D., & de Freitas, N. (2017). *Parallel Multiscale Autoregressive Density Estimation*, DeepMind. Retrieved from http://arxiv.org/abs/1703.03664

Roberts, J. (2016). Thinking Machines: The Search for Artificial Intelligence. *Distillations*. Retrieved from https://www.sciencehistory.org/distillations/magazine/thinking-machines-the-search-for-artificial-intelligence

Rooks, G., Raub, W., Selten, R., & Tazelaar, F. (2000). How Inter-Firm Co-operation Depends Embeddedness : A Vignette Study. *Acta Sociologica*, *43*(2), 123–137.

Ross, S., Moffat, K., McConnachie, A., Gordon, J., & Wilson, P. (1999). Sex and attitude: A randomized vignette study of the management of depression by general practitioners. *British Journal of General Practice*, *49*(438), 17–21.

Sanders, N. R. (2016). How to Use Big Data to Drive Your Supply Chain. *California Management Review*, *58*(3), 26–48. doi:10.1525/cmr.2016.58.3.26

Schechner, S., MacMillan, D., & Lin, L. (2018). U.S. and Chinese Companies Race to Dominate AI. Retrieved from https://www.wsj.com/articles/why-u-s-companies-may-lose-the-ai-race-1516280677

Schield, M. (2004). Information Literacy , Statistical Literacy and Data Literacy. *IASSIST Quarterly*, (6), 6–11. doi:10.1300/J295v03n01{_}22

Schulz-Hardt, S., Frey, D., Lüthgens, C., & Moscovici, S. (2000). Biased information search in group decision making. *Journal of Personality and Social Psychology*, *78*(4), 655–669. Retrieved from http://doi.apa.org/getdoi.cfm?doi=10.1037/0022-3514.78.4.655

Selin, H. (2008). *Encyclopaedia of the History of Science, Technology, and Medicine in Non-Westen Cultures* (2nd). Springer Netherlands.

Shah, S., Horne, A., & Capellá, J. (2012). Good Data Won't Guarantee Good Decisions. *Harvard Business Review*, (April).

Shook, E. & Knickrehm, M. (2017). *Harnessing Revolution: Creating the Future Workforce.* Accenture.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., … Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *Nature*, *550*(7676), 354.

Singh, S. (2011). *The Code Book: The Science of Secrecy from Ancient Egypt to Quantum Cryptography* (Reprint ed). Anchor.

Skitka, L. J., Mosier, K. L., & Burdick, M. (1999). Does automation bias decision-making? *International Journal of Human Computer Studies*, *51*(5), 991–1006. doi:10.1006/ijhc.1999.0252

Skitka, L. J., Mosier, K., & Burdick, M. D. (2000). Accountability and automation bias. *International Journal of Human Computer Studies*, *52*(4), 701–717. doi:10.1006/ijhc.1999.0349

Stratopoulus, B., T & Dhenning. (2000). Does Successfull Investment in Information Technology Solve the Productivity Paradox? *Information & Management*, *38*, 103–117.

Tambe, P., Hitt, L. M., & Brynjolfsson, E. (2012). The Extroverted Firm: How External Information Practices Affect Innovation and Productivity. *Management Science*, *58*(5), 843–859. doi:10.1287/mnsc.1110.1446

The Economist. (2017a). Artificial intelligence will create new kinds of work. *The Economist*. Retrieved from https://www.economist.com/news/business/21727093-humans-will-supply-digital-services-complement-ai-artificial-intelligence-will-create-new

The Economist. (2017b). Automatic For The People. *The Economist*. Retrieved from https://www.economist.com/news/business/21720675-firm-using-algorithm-designed-cern-laboratory-how-germanys-otto-uses

The Economist. (2018). How AI is Spreading Throughout the Supply Chain. *The Economist*. Retrieved from https://www.economist.com/news/special-report/21739428-ai-making-companies-swifter-cleverer-and-leaner-how-ai-spreading-throughout?fsrc=scn/fb/te/bl/ed/howaiisspreadingthroughoutthesupplychaininalgorithmswetrust

The Hackett Group. (2013). Data-Driven Decision Making Leads to More Effective Financial Performance. *Managing Credit, Receivables & Collections*, (July), 4–5.

Thompson, T. (2003). Adherence to advance directives in critical care decision making: vignette study. *Bmj*, *327*(7422), 1011–. doi:10.1136/bmj.327.7422.1011

Turing, A. M. (1936). On computable numbers. *Proceedings of the London Mathematical Society*, *42*, 230–265. doi:10.2307/2268810

Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind*, *59*(236), 433–460.

UCLA: Statistical Consulting Group. (2017). How do I standardize variables in Stata? | Stata FAQ. Retrieved from https://stats.idre.ucla.edu/stata/faq/how-do-i-standardize-variables-in-stata/

Wastyn, A. & Hussinger, K. (2011). In Search for the Not-Invented-Here Syndrome : The Role of Knowledge Sources and Firm Success. In *Druid conference*.

Waters, R. [R.]. (2017). Google Owner Tests First Driverless Car on City Streets. Retrieved from https://www.ft.com/content/397134f8-c36c-11e7-a1d2-6786f39ef675

Weber, J. (1992). Scenarios in business ethics research: Review, critical assessment, and recommendations. *Business Ethics Quarterly*, *2*(2), 137–160.

Weill, P. (1992). The relationship between investment in information technology and firm performance: a study of the value manufacturing sector. *Information Systems Research*, *3*(4), 307–333. doi:10.1287/isre.3.4.307

Wessel, M. (2016). You don't need Big Data — You need the right data. *Harvard Business Review*, 2–6. Retrieved from www.hbr.org

Wickens, C. D., Clegg, B. A., Vieane, A. Z., & Sebok, A. L. (2015). Complacency and Automation Bias in the Use of Imperfect Automation. *Human Factors*, *57*(5), 728–739. doi:10.1177/0018720815581940

Wikipedia. (2018). List of Cognitive Biases. Retrieved from https://en.wikipedia.org/wiki/List_of_cognitive_biases

Appendices

A Dependent Variable Mean Values

Avg Survey 1	Avg Survey 2	EXEC	DATA	Group	N Survey 1	N Survey 2
52.917	51.000	0	0	Control	12	4
68.600	68.889	0	1	Control	10	9
60.778	45.600	1	0	Control	9	5
76.333	71.625	1	1	Control	9	8
41.286	43.875	0	0	Treatment	7	8
84.800	81.600	0	1	Treatment	5	5
76.429	45.333	1	0	Treatment	7	6
62.000	73.600	1	1	Treatment	8	5

Table 12: Scenario 1 Mean Values: HiPPO, Bias and Data

Avg Survey 1	Avg Survey 2	EXT	INT	Group	N Survey 1	N Survey 2
61.778	28.667	0	0	Control	9	3
70.400	57.250	0	1	Control	10	8
31.917	39.333	1	0	Control	12	9
57.000	38.833	1	1	Control	9	6
41.500	44.000	0	0	Treatment	8	5
58.857	50.000	0	1	Treatment	7	2
28.182	40.250	1	0	Treatment	11	12
14.000	34.400	1	1	Treatment	1	5

Table 13: Scenario 2 Mean Values: NIH Syndrome

Avg Survey 1	Avg Survey 2	HUM	MACH	Group	N Survey 1	N Survey 2
71.214	65.857	0	0	Control	14	7
50.467	50.286	0	1	Control	15	7
66.857	59.000	1	0	Control	7	6
46.750	52.833	1	1	Control	4	6
63.000	58.500	0	0	Treatment	7	4
32.500	50.000	0	1	Treatment	6	6
70.333	70.833	1	0	Treatment	6	6
50.875	55.875	1	1	Treatment	8	8

Table 14: Scenario 3 Mean Values: Delegation to Machines