



Universiteit Leiden

ICT in Business

Data Governance practices for Big Data ethics

Name: Nunzia Concetta Bevacqua
Student-no: S2119226

Date: 03/04/2020

1st supervisor: Niels van Weeren
2nd supervisor: Stefan Wolfgang Pickl

MASTER'S THESIS

Leiden Institute of Advanced Computer Science (LIACS)
Leiden University
Niels Bohrweg 1
2333 CA Leiden
The Netherlands

Table of Contents

Abstract	4
1. Introduction	5
1.1. Research Objective	5
1.2. Literature Gap	6
1.3. Research Questions	6
1.4. Relevance	7
1.5. Research Approach	8
2. Literature Review	9
2.1. Conceptual framework	9
2.2. Ethics	10
2.2.1. Business Ethics	12
2.2.2. Ethical Principles	14
2.3. Big Data	19
2.2.1. Big Data supporting technologies	20
2.2.2. Big Data lifecycle	21
2.4. Big Data Ethics	24
2.3.1. Applying ethical values to Big Data	26
2.5. Data laws and regulations	35
2.6. Data Governance	37
2.7. Discussion	40
3. Research Methodology	43
3.1. Conceptual Framework	43
3.2. Methodology approach	44
3.3. Methodology overview	45
3.4. Literature review methodology	46
3.5. Interview methodology	48
3.6. Survey construction	51
3.7. Results validation	53
4. Interview Analysis	55
4.1. Validation of literature framework	55
4.1.1. Respect for Autonomy	55
4.1.2. Nonmaleficence	56
4.1.3. Beneficence	57
4.1.4. Justice	58
4.1.5. Privacy and Data Protection	59
4.1.7. Revised framework	61
4.2. Riskiness level of the Big Data lifecycle phases	62

4.4. Data Governance standard for Big Data ethics.....	63
4.5. Relations between categories	64
5. Survey Analysis.....	66
5.1 Descriptive data.....	66
5.2 Survey results	68
5.2.1. Data collection phase	69
5.2.2. Data cleansing phase.....	71
5.2.3. Data analysis phase	73
5.2.4. Data visualization phase	75
5.2.5. Decision making phase	77
5.2.6. Decision archival and deletion phases.....	80
5.2.7 Risks of unethical Big Data usage	81
5.2.8 The responsibility of the law vs companies	82
5.3. Status of commercial organizations in addressing Big Data ethics	83
5.3.1. Governance practices in use within commercial organizations	84
6. Results Validation	88
6.1. Validation of framework	88
6.1.1. Data Collection.....	88
6.1.2. Data Analysis.....	89
6.1.3. Decision Making.....	91
6.1.4. Data Cleansing	91
6.1.5. Data Visualization.....	91
6.1.6. Data Archival and Data Deletion	92
6.2. Validation of commercial companies' ethical data status.....	92
6.2.1. Frequency of mentions of Data Governance practices.....	92
6.2.1. Perceived importance and priority level of the Big Data lifecycle phases	94
7. Discussion.....	96
7.1. Usability of Data Governance framework for Big Data ethics	96
6.2. Status of commercial companies in addressing Big Data ethics	97
8. Conclusion.....	99
8.1. Research implications	101
8.2. Limitations.....	101
8.3. Future research.....	102
Bibliography.....	103
Annex A.....	108
Annex B.....	109
Annex C.....	110
Annex D	110

Annex E.....	111
Annex F.....	116
Annex G.....	117
Annex H.....	118
Annex I.....	120
Annex J.....	130
Annex K.....	132

Abstract

This study investigates the ethical dilemmas that exist throughout the phases of the Big Data lifecycle, from the moment data is collected, cleansed and analyzed to the moment resulting knowledge is discovered and utilized. Furthermore, it aims to evaluate the ethical preparedness of commercial organizations in collecting personal data, analyzing it and using the outcomes of the analysis to support their decision making purposes, as well as determine how Data Governance practices can aid such organizations in integrating ethics into their data processes. Building on existing work on data ethics, it asks the question: *“How can Data Governance support commercial organizations in addressing Big Data ethics?”*. In this context, Big Data ethics has been defined as a discipline that develops moral standards or practices that support moral decision-making based on Big Data analytics within business environments.

Semi-structured interviews were conducted with 8 experts to validate the findings of the analysis of literature on the concepts of Big Data and ethics. The results showed that commercial organizations engage in unethical activities involving Big Data of personal nature, and that the activities related to the data collection and the data analysis phases of the Big Data lifecycle are particularly impactful for the customer – who is at risk of seeing their human rights and privacy violated. An online survey was then distributed to data practitioners that work at different phases of the Big Data lifecycle. The respondents were asked to make an assessment of which data-related ethical issues they experience within their organizations, and which Data Governance practices they have in place to address them. The analysis of the responses demonstrated that commercial organizations are currently focusing their efforts on reaping the benefits of emerging Big Data analytics technology rather than focusing on handling personal data in respect of the customer’s rights.

A Data Governance standard for Big Data ethics was developed to provide commercial organizations with governance practices they can adopt to address risky, unethical Big Data activities and thus limit the negative impact that unethical data processes can have on the customers. It is recommended that commercial organizations adopt such practices holistically throughout the Big Data lifecycle on the basis of their current level of ethical maturity. Further research is required to translate the standard into a maturity model.

1. Introduction

In recent days, changes in computing power, reduction in data storage costs, better data analysis through Artificial Intelligence, and improvement in networks and the internet have increased the potential for ethical violations within businesses: questions arise about the implications of the acquisition, storage, and use of Big Data (that is, data that is too big to be handled by traditional databases protocols) to infer about people's behavior, preferences, and locations (Wells, 2018). These technological advancements, and related concerns, have paved the way for the definition of a new branch of applied ethics called data ethics, which studies and evaluates moral problems related to data, algorithms and corresponding practices, in order to formulate and support morally good solutions (e.g. right conducts or right values) (Floridi & Taddeo, 2016); it ultimately has the potential of increasing both long term and short term profitability, as well as of developing an increasingly positive public image (Horton, 2019).

Regulations such as GDPR have recently come into effect to give EU residents more control over data, and to ensure that personal data is collected legally and is safeguarded from misuse (Matthews, 2018). While most companies only comply to the GDPR rules in fear of the heavy fines that they could incur in, the Institute of Business Ethics (2018) argues that instead it should be seen as an opportunity to build and sustain customer trust, as well as a way of testing the ethical attitude of an organization. Furthermore, being ethical is not the same as following the law. The law often incorporates ethical standards that most citizens live by; however, it can deviate from what is ethical. This is particularly true when it comes to information technologies and data. The data environment is evolving rapidly (DAMA International, 2017): there are some privacy rules to govern existing flows of personal information, but rules to govern novel flows, uses and decisions derived from that data are missing (Richards & King, 2014).

In the recent years, ethical concerns were raised about the use of information technology to influence the public's opinions, which is not only limited to targeted advertising. In 2016 and 2017, Cambridge Analytica, a data analytics firm, was accused of profiling voters in the United States and elsewhere (Susser, Roessler, & Nissenbaum, 2019), by creating models which used large amounts of Facebook data to influence the voting behaviors of citizens. These algorithms, fed by Big Data, are unnoticeably shaping our lives and decisions (VPRO, 2018). Engaging in such online manipulation practices harms individuals by diminishing their interests and, most importantly, it poses a threat to individual autonomy (Susser, Roessler, & Nissenbaum, 2019). Unfortunately, most organizations lack an ethical culture that prevents such unethical behavior from happening, as well as ethical data governance practices for sourcing and analyzing Big Data. They often invest large amounts of money and effort on developing their analytics capabilities, and yet lack the understanding of how to use them in an ethical manner (Asadi Someh, Breidbach, & Davern, Ethical Implications of Big Data Analytics, 2016). Data Governance can intervene in this context to create an ethical culture, which in practice means introducing controls to ensure that the outcomes of data processing are ethical and do not violate ethical principles, human rights and data regulations (DAMA International, 2017).

1.1. Research Objective

The ethical issues around data raise questions that organizations need to address by translating them into policies and guidelines that help people make ethical decisions in their daily work with data (Wells, 2018). In 2018, the High Level Expert Group on Artificial Intelligence set up by the European Commission proposed a framework achieve and operationalize Trustworthy AI within an organization. The described guidelines, however, apply specifically to the development, deployment and use of AI systems, and does not consider the Big Data lifecycle in its totality. This research wants to identify which ethical issues exist throughout the whole Big Data life cycle, from the moment data is collected, cleansed and analyzed (by means of an AI algorithm or another analytics technique), to the moment resulting knowledge is discovered and utilized. Furthermore, it wants to identify ways to prevent such ethical problems from occurring, by defining Data Governance practices that can support any commercial organization with the task of integrating ethics into their data processes.

These practices will help ensure that an organization can simultaneously manage risk and build trust by consistently evaluating how ethics are taken into account in data-driven decisions. The goal is to have a robust Data Governance program in place that takes data-related ethical issues into account; not only does good Data Governance increase the decision-making ability of those at the top, it also helps everything else run smoothly through an organization (Clark, 2019). The outcomes of this study aim to contribute to the development of successful Data Governance programs within commercial organizations.

1.2. Literature Gap

Existing literature discusses the concept of data ethics, however a formal definition of Big Data ethics has not been provided yet. The existing body of literature did raise ethical questions around the activities involved in turning raw data into insights for decision making purposes: these issues involve, for example, the possible violation of Big Data against user privacy or the practice of collecting public data without seeking appropriate approval (Liu, Li, Li, & Wu, 2015). These activities, however, were never collectively reviewed and formally related to the fundamental ethical principles that govern ethics – which allow to judge in any given field what is ethical and what not.

Ethical questions arise from the things that we do (or don't do) with data – how we collect and use data; the collection and use of data call for ethical judgement. Data ethics is an increasingly important topic of data management and an area where data governance can take a leading role (Wells, 2018). And yet, what emerges from literature is that while Data Governance has been appointed as a way to create an ethical data culture within an organisation (DAMA International, 2017), it was never translated into practices that address unethical data activities occurring within commercial organisations.

When discussing the ethics of data, privacy is often a word associated to it; security also comes up as a related concept because it is what inhibits the unauthorized dissemination of personal data (Wahlstrom, Roddick, Sarre, Estivill-Castro, & deVries, 2006), and thus it supports the protection of user privacy. Alshboul, Wang, & Nepali (2015) suggest, however, that Big Data security should be looked at from different angles and perspectives than the one of security. Furthermore, they identify potential security threats for the privacy of personal data throughout the phases of the Big Data lifecycle. The ethical role of security in the context of Big Data has already been established in scholar literature, and technical security solutions to Big Data-related ethical issues have been proposed. This research will exclude the aspect of security from the review of Big Data ethics and focus on other aspects such as privacy.

1.3. Research Questions

This research wants to identify which unethical activities are performed when using Big Data for decision-making purposes, activities that span throughout the Big Data lifecycle – which describes the cycle of data from the moment it is collected and prepared for its analysis, to the moment that insights are derived and used for decision making.

Furthermore, the research wants to define the role that Data Governance plays in addressing the ethics of Big Data, which in turn will serve to identify Data Governance practices that can support any commercial organization in handling the ethical risks involved in Big Data activities. These practices will constitute the extra step, additional to being compliant to ethical data laws and regulations, that organizations can take to minimize the risks linked with unethical behavior.

The main research question of this research is:

“How can Data Governance support commercial organizations in addressing Big Data ethics?”

In order to answer the main research question, it is first necessary to address the concept of Big Data ethics by providing a definition of it. The sub-question that addresses this task is:

[1] *“What are Big Data ethics?”*

Before diving into Data Governance as a proposed solution to Big Data ethics, the role of data laws and regulations in the debate needs to be addressed, as laws codify some ethical principles but are not equivalent to being ethical. This is discussed in the following sub-question:

[2] *“What is the role of existing data laws and regulations in addressing Big Data ethics?”*

Further, ethical principles are used as a baseline to determine when unethical actions may occur throughout the Big Data lifecycle. By applying these principles to the specific context of Big Data (which is inclusive of the several stages that raw data goes through to be turned into useful insights for decision makers), it should be possible to derive the occurrences of ethical violations when handling Big Data. The following sub-question is proposed to address this matter:

[3] *“How do fundamental ethical principles relate to the Big Data context?”*

Lastly, the possibility of using Data Governance as an instrument for organisations to address unethical Big Data activities will be tackled. This role will be discussed in terms of practices that commercial firms currently adopt and that they should adopt to become more ethical from a data perspective. The following sub-questions are then proposed to address state-of-the-art Data Governance practices and not-yet-adopted, recommended practices:

[4a] *“What Data Governance practices are currently being used by commercial organizations to address Big Data ethics and data regulations?”*

[4b] *“What Data Governance practices should be used by commercial organizations to address Big Data ethics and data regulations?”*

1.4. Relevance

This research represents a first step towards giving importance to the role of Data Governance in addressing data-related ethical issues; this should, in turn, enable organizations to shift their focus from a normal execution of their day-to-day operations with data, to reflecting about how those operations and the decisions derived from them can have a direct or indirect effect on the end user.

Being compliant to regulations is not sufficient to assure a company that they are handling data in an ethical way. As the area of socially responsible and ethical investing keeps growing, companies have more and more of an incentive to be ethical: ethical behaviour is, in fact, increasingly being used by stakeholders (such as investors and customers) to shape their decisions to purchase or invest (Horton, 2019). Unethical behaviour might lead to missing out on these opportunities, as well as incurring into the risk heavy fines and damaging the firm reputation. Firms should be aware that ethical behaviour can bring significant benefits to a business (ACCA, 2014). Besides the fear of consequences, firms should aim to turn ethical behavior into a competitive advantage. By implementing ethical practices, it should be possible to derive business value out of handling data in an ethical way: with pressures coming from new privacy regulations, organizations have a unique opportunity to derive business value out of handling data in an ethical way (Gartner, 2018) – value that goes beyond the simple compliance to ethical data laws and regulations. Ethical data practices can support organizations in the process of evaluating how ethics influence their data-driven decisions; by focusing on ethics, organizations can improve the trust their customers have in them (Tiell & O'Connor, 2016). Also, being more considerate of the ethical implications of working with data can help an organization make quicker and better decisions, as well as support it to being compliant (Clark, 2019).

Organisations that have the ability of unlocking value from their data faster than their competitors will likely be the winners in the race to see who can get the most benefit from Big Data, and Data Governance practices will likely reveal themselves to be instrumental in this race (Tallon, 2013). Ethics will become the new parameter for competitive advantage, and only companies with the highest of morals and a governance framework to support it will succeed (Knudsen, 2019). This research wants to contribute to the development of concrete measures, in the form of Data Governance practices, that firms can adopt to support their journey towards a more ethical handling of personal data.

1.5. Research Approach

This research will investigate the Big Data context in order to identify unethical activities, and then point out Data Governance practices that organizations can take into consideration to make a more ethical use of their data. The research will attempt to reach its objectives by applying of a combination of qualitative and quantitative methods that will both will work towards filling the identified literature gap.

- 1) The literature review will serve to define the concept of Big Data ethics, building it up through the definition of Ethics and Business Ethics. The role of laws and regulations in addressing Big Data ethics will be addressed. Furthermore, the concept of ethics will be operationalized in order to identify risky, unethical activities involving the use of Big Data.
- 2) Expert interviews will be used to validate the findings of the literature review and to investigate Data Governance practices that may be used by commercial organizations to mitigate the risks of handling data unethically.
- 3) A survey addressing data practitioners will serve to assess how prepared commercial organizations are in addressing Big Data ethics, as well as to understand what practices it would be desirable for them to use to tackle unethical Big Data activities they might engage in.
- 4) A validation session with a group of experts will be used to validate the data collected in the previous stages of the research.

The research process is summarized in the figure below:



FIGURE 1 – OVERVIEW OF THE RESEARCH APPROACH

The deliverable of this research will be a Data Governance standard for Big Data ethics which provides its user with a list of unethical Big Data activities that carry a certain level of ethical risk, and a control for each of them in the form of Data Governance practices. The standard could be used to govern Big Data within a commercial firm and transform their data process into one that takes ethics into consideration – thus ensuring an ethical handling of Big Data within the organization itself. The framework will result from the qualitative data collected during the interviews with experts – which will provide the foundational structure of the framework, and survey with data practitioners – which will provide new information to add to the framework. Lastly, the framework will be validated with a panel of experts – which will generate the final draft of the standard.

2. Literature Review

A theoretical framework is a guiding structure for a research, and consists of a theory constructed by using an established explanation of certain phenomena and relationships. Such theory underlies and supports the research plan, and it also provides a basis of relevant concepts and definitions that will be used for understanding and analysing the research topic. For qualitative research, which has a more exploratory nature, the theoretical framework may be less structured to keep the researcher from forcing preconceptions on the findings (Grant & Osanloo, 2014).

On the other hand, a conceptual framework describes the relationship between the main concepts of a study, and it usually visually displays how ideas in the study relate to one another (Adom, Joe, & Hussein, 2018). A conceptual framework allows the researcher to specify and define concepts within the problem. It provides insights on how the researcher will explore the research problem, which direction the research will take, and the relationships between the variables in the study (Grant & Osanloo, 2014).

In this chapter, a definition of the core concepts of the research topic will be provided, such concepts being *ethics, business ethics, Big Data and Big Data ethics, and Data Governance*.

2.1. Conceptual framework

When discussing Big Data, we can think of the series of activities associated to it: these involve the collection of data, its integration and cleaning, its analysis supported by a specific technology such as ML and lastly the informed decision making. Floridi & Taddeo (2016) argue that ethical problems such as privacy, transparency, trust and responsibility concern the lifecycle phases of data collection, curation, analysis and use, and hence they are better understood when analysed at this level. Ethical principles serve to define ground rules for ethics, and they help us distinguish between what is ethical and what not. Such principles, once identified, can be applied to the specific context of Big Data to understand where ethical violations may occur within the cycle of Big Data activities.

Laws and regulations also play a role in the data ethics debate because they are supposed to be there to regulate Big Data activities and prevent unethical ones to occur in the first place. However, while laws are usually founded on the same ethical principles discussed in the previous paragraph, being compliant to them often does not correspond to being ethical. There have been instances in which firms were able to use Big Data in unethical ways, while technically not breaking the law. The relationship between the Big Data activities and laws and regulations closes with a feedback loop: that is because sometimes laws do not take certain unethical activities into account until they occur – and this is especially true when new technologies arise and new ethical problems arise with them. Thus, there is the need for governments to be aware of unregulated unethical activities that happen when using Big Data and update their laws to take them into account. But before that happens, what are organisations supposed to do?

Ideally, unethical Big Data activities should be identified and tackled without having to first wait for laws to be updated and regulate them. The hypothesis for this research is that Data Governance can intervene to support the design of ethical Big Data systems. Data Governance should have the goal of reducing the risk of improper behaviour by setting appropriate ethical standards and policies; by taking people, processes and technology and oversee them, Data Governance could prevent unethical behaviours from occurring .

The relationships between the major concepts of the research are depicted in the figure below:

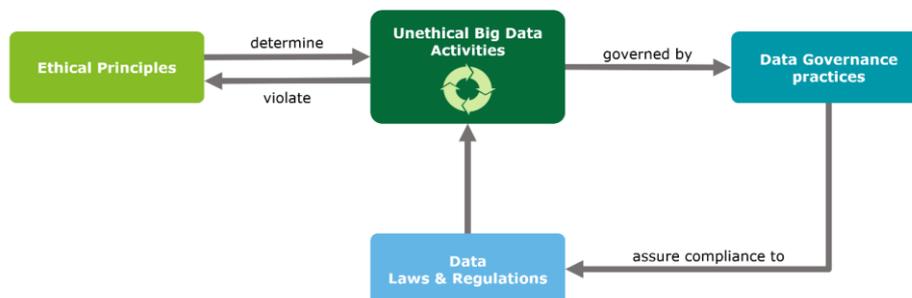


FIGURE 2 – RELATIONSHIPS BETWEEN MAJOR RESEARCH CONCEPTS

2.2. Ethics

The term *ethics* is often used interchangeably with the term *morality*. The words derive respectively from the Greek *ethos* and the Latin *mores*, both of which refer to the notions of custom, habit, behaviour and character; however, there is a distinction to be made between the two words.

Morality

Gert (1999) provides the following definition of morality:

“Morality is an informal public system applying to all rational persons, governing behaviour that affects others, and includes what are commonly known as the moral rules, ideals, and virtues and has the lessening of evil or harm as its goal.”

Morality is described as a system which applies to all rational persons who are responsible for their actions. All rational persons are moral agents, bound by the system of moral rules. This system is informal because it has no formal authoritative judges presiding over it. Morality is also public because everyone must know what the rules that define it are, and every moral agent is obligated to participate in it. Morality is comprised of moral rules – which prescribe what humans ought to do in terms of obligations (e.g. refrain from stealing and murdering), virtues (e.g. honesty, compassion, and loyalty), rights (e.g. the right to privacy), and values – which are ends or goals sought by individuals (e.g. health and happiness). The purpose of morality is to prevent harms and evils.

(Tavani, 2004) (Gert, 1999) (Velasquez, Adre, Shanks, J., & Meyer, 2010)

Tavani (2004) defines morality as a system comprised of rules guiding human conduct – which describe what people ought and ought not to do, and principles for evaluating those rules. Moral rules are rules of conduct that can take the form of:

1. Directives that guide our conduct as individuals at a microlevel – the level of individual behaviour. An example of these is “Do not harm others”.
2. Social policies framed at a macrolevel – the level of social policies and social norms. An example of these is “Software that can be used to invade the privacy of users should not be developed”.

The rules of conduct in a moral system are evaluated against standards called principles. Such principles, or morals, are standards of behaviour which are used to determine what is right and wrong, and which can be used for determining whether policies can be justified on moral grounds.

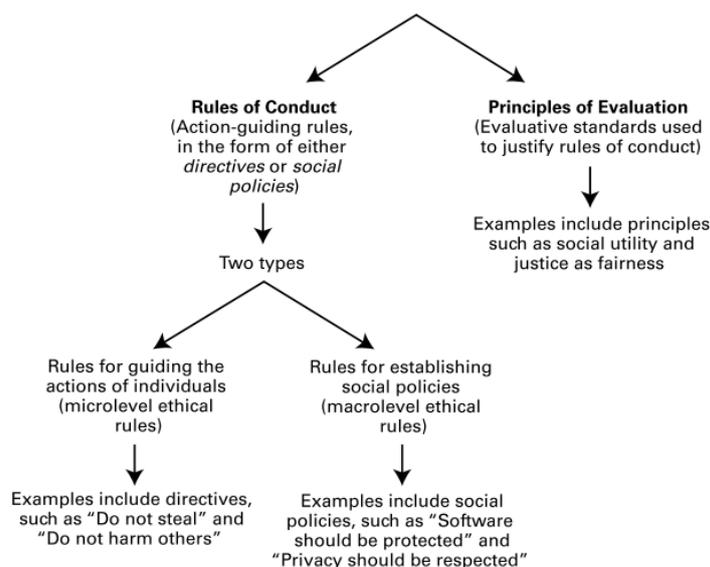


FIGURE 3 - DEFINITION OF MORALITY (TAVANI, 2004)

Since sources are not consistent when describing morality, this research proposes an alternative definition of it. This highlights the informal nature of morality – whose rules are not formally written on a Constitution of Charter, the elements that it is constituted of, and its goal.

“Morality is an informal system comprised of virtues, rights, values, rules, as well as principles that evaluate such rules, which aims to prevent harms and evils”.

Ethics

Ethics is defined as the philosophical study of morality, which examines people’s moral beliefs and behaviour (Quinn, 2004) (Tavani, 2004). Ethics:

- Study the moral choices that people make and the way in which they seek to justify them.
- Develop ethical standards, ensuring that these are reasonable and well-founded.
- Ensures that humans and institutions live up to reasonable and solid ethical standards.

Ethics deals with moral principles, and thus with distinguishing between good and bad judgements, right and wrong and what ought to be. Ethics imply the examination of morality in the context of individual and social behaviour, religion, culture, and personal life. People exercise ethics when they face predicament or problems concerned with morality; in these situations, they require a framework of ethical principles and values that can support them in resolving these problems. Ethical theories investigate morality and provide such principles and values so that individuals and society are able to make decisions when facing moral predicaments or problems. (Velasquez, Adre, Shanks, J., & Meyer, 2010) (Singh & Mishra, 2018)

In this research, an alternative definition of ethics is proposed to gather the opinions of multiple sources:

“Ethics is a discipline concerned with the moral choices of humans, which develops reasonable ethical standards that support moral decision-making and ensures that society lives up to such moral principles.”

This definition highlights the multiple perspectives of ethics as a field of study: a philosophical (normative) perspective concerned with creating moral standards, and a behavioural (descriptive) perspective concerned with describing how people behave and make moral choices.

2.1.1. Business Ethics

Amakobe (2016) provides the following definition of business ethics:

“Business ethics is a form of applied ethics that examines ethical principles and moral or ethical problems that arise in the business environment. It applies to all aspects of business conduct. It is relevant to the conduct of individuals and business organizations as a whole.”

Business ethics is the study of morality in the sphere of business organisations, and can be thought of as an area of applied ethics. There are two possible approaches to considering ethics:

1. Normative: it involves creating moral standards, which means deciding what is ethically correct to do. Facing a decision, it determines what is the “right” course of action. This approach prescribes what one ought to do and what guiding values one should embrace. It is also concerned with evaluating people’s moral behaviour and determining whether it’s reasonable.
2. Descriptive: it is concerned with describing how people behave and understanding what makes them behave in a certain way when confronted with moral choices (for example by describing what moral standards they follow). Differently from the normative approach, the descriptive approach attempts to describe what is, and not what ought to be.

(Amakobe, 2016) (Wittmer, 2009)

Some authors prefer to provide a definition of business ethics that focuses solely on one of the possible approaches of business ethics, either normative or descriptive. For example, Fischer (2004) has a descriptive view of business ethics and states that:

“Business ethics is considered with the actual customs, attitudes, values and mores that operate within business.”

On the other hand, Lewis (1985) provides a definition of business ethics that solely considers the normative approach:

“Business ethics is rules, standards, codes or principles which provide guidelines for morally right behaviour and truthfulness in specific situations”.

In this definition, business ethics is described as a discipline that provides moral guidelines that, if followed, will prevent unethical behaviour; it provides the standards or principles that guide behaviour, specifically in the context of business. Fischer (2004) argues that the rightfulness of a specific behaviour is often determined by stakeholders, such as investors, customers, interest groups, employees, the legal system, and the community. What’s ethical or unethical is determined by reference to what is considered acceptable behaviour for business. Lewis (1985) similarly states that behaving morally right means that actions confirm to justice, law, or other standards. According to Lewis it also means being aware of the consequences of one’s actions and refusing to engage in practices that would corrupt one’s integrity. Business ethics involves applying one’s understanding of what is morally right and truthful at a time of moral dilemma.

(Fischer, 2004) (Lewis, 1985)

The general goal of business ethics is to guide and support individual decision making and to develop an ethical workplace environment (Fischer, 2004). It does so, as explained above, by preventing ethical problems and unethical behaviour from occurring within a business. Business ethics has the functions of: (1) helping business people make policies and strategies to avoid major frauds and scandals and to make the business ethical overall; (2) supporting the understanding of the very foundation of business, which is rooted in human morality and social ethics; (3) providing a framework of rules, principles, and values. Furthermore, it works to protect the business from big damages, detect unethical behaviour and develop ethical strategies that could eliminate unethical behaviour and practices within a company. Business ethics works internally within business organisations to alert, avoid and eliminate the malpractices, mismanagement and wrongdoings; it also works externally to put pressure in the context of society so that unethical malpractices could be eliminated from organisations and the values of society, the environment and stakeholders could be protected.

(Singh & Mishra, 2018)

From the existing literature, we can derive a definition of business ethics which is inclusive of the different perspective of the scholars who discussed this topic. The following definition is proposed:

“Business ethics is the study of morality within the business environment. Its goal is to develop an ethical workplace environment, and it does so by both: (1) providing rules, moral principles and values; (2) describing and explaining the behaviour of business people in situations of moral dilemma.”

This definition provides an explanation of what business ethics is, clarifies its goal and describes its functions, while remaining inclusive of both the normative and descriptive perspectives.

Importance of business ethics

De George (1987) argues that the importance of business ethics lies in its potential to raise ethical issues within a firm. Ethics will not solve business problems nor replace any area of business education. It does, however, look at the ethical implications of one’s actions and helps one become more objective. It also has the potential of preventing disasters from happening.

The importance of business ethics lies in its relation with a company’s reputation, which is one of a company’s most important assets: easy to lose, and difficult to rebuild. Businesses that do not follow any kind of ethical code and behave unethically risk damaging their own reputation and looking less appealing to investors and customers. Profits could fall as a result of this.

Ethical behaviour shouldn’t, however, only be driven by the fear of consequences. Exercising corporate social responsibility can bring significant benefits to a business. It may help attract more customers – and thus boost sales, retain employees – and consequently increase productivity, as well as attract new employees and investors. From the point of view of the employee, working for a company with strong business ethics reassures them that the firm will not allow unethical practices to occur; from the point of view of the customer, knowing that they are buying products from an ethical and responsible company makes them feel more ease. Just as the element of sustainability can push people to buy a specific brand of coffee, ethics could be integrated in a product strategy and become a selling point for the final product.

(ACCA, 2014) (Horton, 2019)

Ultimately, setting up and following ethical guidelines can reduce the risk of a firm of being fined for poor behaviour due to issues of non-compliance to law. The more customers and investors seek to purchase products and invest in an ethically operating company, the more incentives for being socially responsible will grow, leading firms to take ethical issues more seriously.

Thus, corporate leaders should strongly consider understanding and applying ethical principles to their day-to-day operations. Unethical behaviour should not be tolerated in any organisation; it should not be justified either because ethics are well understood, and some of them may even be codified into procedures, processes, rules, regulations and laws (Amakobe, 2016).

Business ethics and regulations

Regulations exist to control and forbid unethical business practices, so that governments don’t have to rely on the independent ethical choices of individuals; they also assure each firm that its competitors won’t be able to get ahead by taking the “low road”. However, according to Norman (2013) businesses often try to resist the imposition of regulations on their activities. Furthermore, regulations will never eliminate all unethical business practices. This is because some regulations would violate basic liberties, some would bring more costs than benefits, some would be too difficult to monitor effectively, and some would be too slow to deal with problems raised by technological innovations.

When the concept of corporate social responsibility came to be known as the moral dimension of business in the mid-nineties, its advocates argued that ethical management requires more than confirming to the law; ethical management should instead anticipate the law by voluntarily undertaking socially responsible actions

that go beyond the minimum legal requirements. They argued that behaving ethically not only would prevent the consequences of non-compliance to regulations, but that the market would ultimately reward such behaviour (Stark, 1993). While the concept of corporate social responsibility has been surpassed by the concept of 'business ethics', the argument remains that business value could be derived by exercising ethical behaviour that goes beyond the simple compliance.

2.1.2. Ethical Principles

It has been discussed that business ethics, unlike the ethics introduced in the previous sub-chapter, apply to particular, concrete situations in the context of business. Based on the provided definitions of ethics and business ethics, it is possible to observe that they both have within their objectives the development of moral principles or standards that guide human behaviour. The drivers of behaviour, however, differ based on the context the moral dilemma presents itself; firms and individuals have different interests and drivers, and different legal and social obligations they have to abide by, thus the consequences of their decisions also differ.

However, ethics and business ethics have something in common. The principles of ethics that help us solve ethical dilemmas in everyday life are the same principles that provide guidance in business, health care, law, and education (Weinstein, 2017). Ethical values such as honesty, transparency, fairness, accountability and integrity apply to both businesses and individuals. Rather than the values themselves, what differs between ethics and business ethics is how the values are applied to specific situations.

(Kramer, 2019)

Values are qualities that one should show and exert in the way they behave. They are distinct from principles, which are instead rules established on values, that govern one's behaviour (Pediaa, 2016). However, the authors mentioned in this sub-section tend to use the words *principle* and *value* interchangeably. The researcher in this paper will seldom take on the same convention, while being aware of the difference in significance of the two terms.

In this section, fundamental ethical values will be introduced: these constitute the foundation of an ethical study. While usually ethical values are discussed in the context of a specific field, the argument is that they are universally applicable to any situation.

Beauchamp & Childress (2001) introduce four ethical principles and show how these apply to biomedical sciences. The same principles are then discussed by Wright (2011) in relation to information technologies, arguing that such principles should be accounted for when executing an ethical impact assessment of a specific technology. Under these major principles, Wright mentions one or more ethical value or issue that have a relation with each overarching principle.

The first four principles in this section, namely *Respect for Autonomy*, *Nonmaleficence*, *Beneficence* and *Justice* are discussed by both Beauchamp & Childress (2001) and Wright (2011). However, Wright also includes the principle of *Privacy and Data Protection* when discussing ethics applied specifically to information technologies. This principle is included in the review due to the nature of this research that fits with what the principle represents.

Respect for autonomy

Personal autonomy is a value that Beauchamp & Childress (2001) define as:

“Personal autonomy is, [...], self-rule that is free from both controlling interference by others and from limitations, [...], that prevent meaningful choice. A person of diminished autonomy, by contrast, is in some respects controlled by others or incapable of deliberating or acting on the basis of his or her desire and plans”.

An autonomous individual freely acts in accordance to their chosen plan; in contrast, a person with no autonomy is partly or fully controlled by others or incapable of acting on the basis of their desires and plans. Violating a person's autonomy means treating them as a means – as an instrument to pursue other's goals, disregarding that person's own goals. Such behaviour is a fundamental moral violation because it constraints a person from shaping his or her own life. Autonomy, equated with liberty, is a right enshrined in Article 6 of the

European Charter of Fundamental Rights as well as Article 3 of the UN's Universal Declaration of Human Rights of 10 December 1948.

Respect for autonomy is an ethical principle derived from the value of personal autonomy, which involves acting in such a way that other persons are unable to act autonomously; consequently, disrespect for autonomy means demeaning other's autonomy. Such principle is stated as: "Autonomous actions should not be subjected to controlling constraints by others".

(Beauchamp & Childress, 2001) (Wright, 2011)

Wright (2011) describes the values and issues of *Dignity*, *Informed Consent* and *Social Solidarity*, which are strictly related to the value of Autonomy. These are listed below:

- **Dignity:** "Citizens should be treated fairly regardless of age, gender, racial or ethnic background, disability or other status, and be valued independently of their economic contribution"; "all human beings are born free and equal in dignity and rights".
Dignity is a right enshrined in Article 1 of the Charter of Fundamental Rights as well as Article 1 of the UN's Universal Declaration of Human Rights.
- **Informed consent:** "[...] personal data may be processed only if: (a) the data subject has unambiguously given his consent". For example, online services should obtain informed consent prior to the collection and use of personal data.
Dignity is a right enshrined in the EU Directive on clinical trials (2001/20/EC) as well as Article 7 of the EU Data Protection Directive.
- **Social solidarity, inclusion and exclusion:** "E-Inclusion refers to the actions to realise an inclusive information society, that is, an information society for all". Cost and knowledge are among the prime reasons why some people are excluded from the information society.
The concept of isolation is mentioned in the European Council resolution (2001/C 292/02) on e-Inclusion.

Nonmaleficence

The principle of nonmaleficence is defined by Beauchamp & Childress (2011) as follows:

"Nonmaleficence is the obligation to not inflict harm intentionally. In contrast with the principle of beneficence, which requires helping, preventing, removing harm and promoting good, nonmaleficence only requires the intentional refrain from actions that cause harm".

In cases of conflict, nonmaleficence typically overrides beneficence, although the way these moral principles apply is highly dependent on the context; there is no a priori rule that favours avoiding harm (nonmaleficence) over providing benefit (beneficence).

(Beauchamp & Childress, 2001)

Wright (2011) describes the values and issues of *Safety*, *Isolation and substitution of human contact*, *Discrimination and social sorting* in relation to the principle of nonmaleficence. These are listed below:

- **Safety:** "In order to [...] ensure a high level of consumer protection, the Community shall contribute to protecting the health, safety and economic interests of consumers, as well as to promoting their right to information, education and to organise themselves in order to safeguard their interests".
Consumer protection is provided by Article 38 of the Charter of Fundamental Rights as well as Article 153 of the EC Treaty, and at European level by Directive 93/13, Directive 97/7 and Directive (85/374/EEC).
- **Isolation and substitution of human contact:** "Isolation is the objective condition of having too few and too poor social ties, of not being in any relevant social network. new communication tools may become a substitution for face-to-face contact and could, thereby, make social isolation worse".

Isolation is a potential issue raised by emerging technologies.

- **Discrimination and social sorting:** “Any discrimination based on any ground such as sex, race, colour, ethnic or social origin, genetic features, language, religion or belief, political or any other opinion, membership of a national minority, property, birth, disability, age or sexual orientation is prohibited”.

Discrimination is a prohibited by Article 21 of the European Charter of Fundamental Rights as well as the Directive adopted by the European Parliament on April 2nd 2009.

Beneficence

The principle of beneficence defined in Iserson (1999)’s paper as:

“Beneficence is the duty to help others further their important and legitimate interests”.

According to Beauchamp & Childress (2001) morality does not only require us to refrain from harming others and allowing them to decide autonomously for themselves; it also requires us to actively contribute to their welfare. The principle of beneficence is thus more demanding than the before mentioned principle of nonmaleficence because it expects a moral agent to not only refrain from harmful acts, but to take action to help others.

Wright (2011) mentions the values and issues of *Universal Service*, *Accessibility*, *Value sensitive design* and *Sustainability* when discussing the principle of Beneficence. These are the following:

- **Universal service:** “Universal service is an obligation imposed on one or more operators of electronic communications networks and/or services to provide a minimum set of services to all users, regardless of their geographical location within the national territory, at an affordable price”. The right to Universal service is enshrined in the EU Directive (2002/22/EC) on universal service and users’ rights relating to electronic communications.
- **Accessibility:** accessibility refers to the user-friendliness of devices and services and is a “prerequisite for the e-inclusion of citizens in the Information Society”. The European Commission has developed an action plan in 2007 to “both to help older people towards a safer and more independent old age and to promote the development of Information and Communication Technologies (ICTs) in services to persons” (European Commission, 2007).
- **Value sensitive design:** “the value of members of a design team, [...], often shape a project in significant ways [...]. Beliefs and commitments, and ethnic, economic, and disciplinary training and education, may frame their perspectives, preferences, and design tendencies, resulting eventually in features that affect the values embodied in particular systems”. The importance of value sensitive designs has been highlighted in a report of the SoBigData project founded by the EU, which states that it would be helpful if digital products and services could in some way send honest signals to users about their moral quality and the values that have been used to shape them, in order to achieve transparency and accountability (Hänold, et al., 2016).
- **Sustainability:** “Sustainability refers to a condition whereby a project or service can be sustained, can continue into the future, either because it can generate the financial return necessary for doing so or that it has external support (e.g., government funding) which is not likely to go away in the foreseeable future”.

Justice

The principle of justice concerns the distribution of social benefits and burdens (Iserson, 1999). The principle is formulated by Wright (2011) as follows:

“Justice is fair, equitable, and appropriate treatment in light of what is due or owed to persons. An injustice thus involves a wrongful act or omission that denies people benefits to which they have a right or distributes burdens unfairly.”

Justice requires from a moral actor that others are equals are treated equally, and the unequal are treated unequally, in proportion to their relevant inequalities (Iserson, 1999). Justice implies in a considerate way towards other people’s interest, property and safety (Wright, 2011).

Wright (2011) associates the concept of justice to the terms ‘equality’ and ‘fairness’, stating that they all appeal to the idea of giving people what they deserve and ensuring that people receive their fair share of benefits and burdens – thus making them all equal.

Privacy and Data Protection

Privacy refers to the right of an individual to have control over the access of his or her personal information. Brian Dickson, a Canadian judge, in the court case *R v Duarte* defines privacy as “the right of the individual to determine when, how, and to what extent he or she will release personal information”. He continues saying that “an individual may proceed on the assumption that the state may only violate this right [...] when it has established [...] that an offence has been or is being committed [...].”

Privacy is a right enshrined in: Article 12 of the Universal Declaration of Human Rights, which states that “No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence”; the European Charter of Fundamental Rights, which states that “EU citizens have the right to protection to their personal data”; the European Convention of Human Rights. In 2016 the General Data Protection Regulation (GDPR) was established to protect the data privacy rights of individuals.

Wright (2011) discusses the issues of *Collection limitation and retention*, *Data quality*, *Purpose specification*, *Use limitation*, *Confidentiality*, *security and protection of data*, *Transparency*, *Individual participation and access to data*, and *Anonymity*. These issues are listed below:

- **Collection limitation and retention:** “There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject”. Data limitation is mentioned in the OECD guidelines, which are recommendations for enterprises addressed by governments in the form of principles and standards for responsible business conduct.
Data retention refers to the storage of data in the form of phone records, transactions, internet information such as emails sent and received and websites visited. Data retention is discussed in the Article 17 of the GDPR, which entitles a data owner to demand erasure of their personal data from a database.
- **Data quality:** “Personal data must be accurate and, where necessary, kept up to date”. Data quality is referred to by GDPR as ‘data accuracy’ in the Article 16, which refers not only to the right of a user to correct inaccurate or incomplete data, but also to the responsibility of organisations to ensure the accuracy of data collected from data subjects.
- **Purpose specification:** “Personal data must be collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes”; “the purposes for which personal data are collected should be specified not later than at the time of data collection”. Purpose specification is mentioned in the OECD guidelines but is also imposed by Article 13 of the GDPR, which imposes the data collector to provide the data subject with the purpose of the processing of their personal data.
- **Use limitation:** “Personal data should not be disclosed, made available or otherwise used for purposes other than those specified except with the consent of the data subject or by the authority of law”. Purpose limitation is mentioned in the OECD guidelines and is also a requirement of Article 6 of the GDPR, which states that personal data collected for one purpose should not be used for a new, incompatible purpose. According to Wright (2011) this principle also refers to the migration of data to

other sources other than the original collector.

- **Confidentiality, security and protection of data:** “Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorised access, destruction, use, modification or disclosure of data.”
Protection of personal data is mentioned in the OECD guidelines, as well as in the GDPR; while the latter does not define exactly the security measures that an organisation is expected to have, it does enforce that the data controller has to implement “appropriate technical and organisational measures to ensure a level of security appropriate to the risk”.
- **Transparency:** transparency refers to being open about the nature of personal data, as well as the main purpose of its use.
The GDPR enforces through Article 12 that organisations are transparent about the way their process user data. Communications should be in a “concise, transparent, intelligible and easily accessible form”.
- **Individual participation and access to data:** “An individual should have the right to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to him, [...], and to challenge data relating to him and, if the challenge is successful to have the data erased, rectified, completed or amended”.
The GDPR enforces the right of individuals to participate in the data process through Article 15 (right of access by the data subject to information concerning, for example, the purpose of data processing and the categories of personal data concerned), Article 16 (right to rectification of data when incorrect and/or incomplete) and Article 17 (right to erasure personal data on the data subject’s request).
- **Anonymity:** “Anonymity ensures that a subject may use a resource or service without disclosing his or her identity”.
This issue is discussed in the ISO/IEC 15409 standard on evaluation criteria for IT security, as well as in Article 9 and 26 of the EU Directive 2002/58/EC on privacy and electronic communications, stating respectively that anonymous data should be used where possible, and that data should be made anonymous after the provision of a service.

Wright discusses privacy by referencing Roger Clarke, who identified four dimensions to privacy: *Privacy of personal communications*, *Privacy of the person*, *Privacy of personal behaviour*, and *Privacy of personal data*. Wright only discusses the first three kinds without motivating why a separate paragraph isn’t dedicated to Privacy of personal data.

Clarke (2016) talks about privacy of personal data saying that personal data “should not be automatically available to other individuals and organisations, and that, even where data is possessed by other party, the individual must be able to exercise a substantial degree of control over that data and its use” (Clarke, 2016). Arguably, exercising control over one’s personal data means being in control of its quality, of its collection and retention, of the use that is made with it; these problematics have already been partly addressed by the issues listed above. Below, the remainder privacy dimensions are listed:

- **Privacy of personal communications:** this aspect of privacy refers to the confidentiality of communications. The listening, tapping and storage of any kinds of interceptions or surveillance of communications by persons other than the user himself should be prohibited (Wright, 2011).
- **Privacy of the person:** this aspect of privacy is concerned with the integrity of the individual’s body. An example of bodily privacy is blood transfusion without consent (Clarke, 2016).

- **Privacy of personal behaviour:** this dimension of privacy refers to all aspects of behaviour, but especially to sensitive matters such as sexual preferences, political activities and religious practices (Clarke, 2016).

2.3. Big Data

Today we are generating 2.5 quintillion bytes of data, which comes from a multitude of sources such as sensors, social media sites, digital pictures and videos, transaction records, location services and smartphones. This abundance of data is referred to as 'Big Data'. "Big Data is about capturing, storing, sharing, evaluating and acting upon information that humans and devices create and distribute using computer-based technologies and networks" (Herschel & Miori, 2017). The rise of Big Data was facilitated by the rapidly expanding, cheaper and highly networked computing capacity (Metcalf, Keller, & Boyd, 2016).

Data carries information that, once it is made sense of, can help organisations make informed decisions and potentially provide competitive advantage. However, due to the fact that the amount of data generated surpassed the capabilities of existing data storage techniques, the current challenge is to be able to store such quantities of data, as well as to analyse it in a timely manner (Bhadani & Jothimani, 2016).

When the concept of Big Data arose, many definitions referred to it as datasets whose size is beyond the ability of general computers to capture, store, manage and process. Size, however, is not the only defining characteristic of Big Data, which led subsequent definitions to define further characteristics. The 3V model first arose to describe the characteristics of Volume, Veracity and Variety (Bhadani & Jothimani, 2016). These are described below:

- 1) *Volume:* the rise of new sources of data generation have accelerated the growth of data volume, which is projected to be 40 zettabytes by 2020 (Herschel & Miori, 2017). Volume thus refers to the magnitude of data that is being generated and collected (Bhadani & Jothimani, 2016), which surpasses the capabilities of traditional storing and analysis techniques (Sagiroglu & Sinanc, 2013).
- 2) *Velocity:* improvements to the telecommunications infrastructure, as well as the deployment of high-speed wireless technologies have increased the velocity of data, that is the speed at which data is transferred and shared globally (Herschel & Miori, 2017). Velocity can also refer to a requirement of processes built around Big Data: Big Data should be analysed in real-time as it streams into an organisation to make informed decisions and in order to maximise its value (Sagiroglu & Sinanc, 2013) (Bhadani & Jothimani, 2016).
- 3) *Variety (or complexity):* it refers to the types of data that are being generated and captured, and the increasing range of formats and representations employed (Ward & Barker, 2013), which makes it difficult to use data collected across different systems (Herschel & Miori, 2017). Data can be structured, semi-structured and unstructured. Structured data refers to data that can be organised using a pre-defined data model; this only constitutes 5% of existing data (such as data in an Excel file). Unstructured data is, by contrast, data that cannot be organised using these pre-defined models (such as video, text and audio). Semi-structured data lies in between the previous two categories (an example is Extensible Markup Language, or XML, data) (Bhadani & Jothimani, 2016).

In addition to the above described qualities of Big Data, later on the dimensions of Veracity and Variability were added to its definition as are referred to as 5Vs. Bhadani & Jothimani (2016) also describe the quality of Low-value density. These dimensions are described below:

- 4) *Veracity:* it refers to the unreliability associated with the data sources, which means that the data being collected and shared might be incomplete and/or inaccurate (Herschel & Miori, 2017). This dimension raises questions of trust and uncertainty with regards to data and the outcome of analysis of that data (Ward & Barker, 2013). It thus brings up the need and challenge to separate reliable data from imprecise data, and to manage such uncertainty. (Bhadani & Jothimani, 2016)

- 5) *Variability*: it refers to the inconsistency of data flows, which can have periodic peaks (Herschel & Miori, 2017). The variation in flow rate of data is often caused by inconsistencies in the velocity of data. (Bhadani & Jothimani, 2016)
- 6) *Low-value density*: this refers to the fact that data in its original form is unusable, and that it must be analysed to derive value from it. (Bhadani & Jothimani, 2016)

De Mauro, Greco & Grimaldi (2015) also analysed existing literature to work towards an inclusive definition of Big Data. They notice the following core concepts in research on Big Data do not only lie in the characteristics described above, but also in specific technology and analytical methods which are a needed requirement to make use of data in the first place, as well as in the transformation into insights which is the way Big Data can impact companies and society through the creation of economic Value. Based on these core concepts, they propose the following definition:

“Big Data represents the Information assets characterized by such a High Volume, Velocity and Variety to require specific Technology and Analytical Methods for its transformation into Value.”

The definitions collected so far are all deemed valid in describing Big Data. While some preferred to focus on the objective characteristics of this phenomenon – in the form of 3Vs or 5Vs, others integrated in their definition the challenges that it brings (such as the required capabilities to collect, store and analyse such big amounts of data). Some others described Big Data as not only a technological phenomenon, but also cultural, referring to the ability to derive insights from this data, as then value from these insights. It is therefore being recognised that Big Data holds immense social and economic value, due to the ability to capture knowledge from data, and to act upon the generated knowledge. Big data boosts the economy by creating new opportunities through the use of analytics; it advances scientific research by opening it up to data-driven discoveries; it supports nations in optimising natural resources, responding to national disasters and enhancing critical information infrastructure (Tene & Polonetsky, 2013). Big Data made us rethink how knowledge can be generated, how research can be executed, how information can be engaged with (Boyd & Crawford, 2012). However, the societal benefits of Big Data must be faced off against the increased risks it brings – such as to individuals’ privacy. (Tene & Polonetsky, 2013)

2.2.1. Big Data supporting technologies

In Microsoft’s definition of Big Data, attention is given to the process used to process complex sets of data. In fact, the term Big Data is often associated with the specific technologies that enable its utilisation. While originally statistical techniques were used to process data, these have been surpassed by more advanced processing methods that are able to analyse the extensive quantities of data available today. Artificial Intelligence and Machine learning are often mentioned as related technologies, and constitute a crucial part of the definition of Big Data (Ward & Barker, 2013) (De Mauro, Greco, & Grimaldi, 2015).

The researcher agrees on this definition and finds important to mention two of the most mentioned technologies in literature that allow to turn raw Big Data into useful insights, namely Artificial Intelligence/Machine Learning and Data Mining (a technology involving methods that are at the interception between machine learning and more traditional statistical methods).

Data Mining

Data Mining is defined as the application of specific algorithms to extract new information from existing data by identifying patterns, correlations or trends in specific categories from the data (Tavani, 2004) (Cary, Wen, & Mahatanankoon, 2003). It is sometimes referred to as Knowledge Discovery because data miners do not exactly know what they are looking for before they find it. The goal of data miners discover new insights from the data in their databases (Cary, Wen, & Mahatanankoon, 2003): the discovered patterns represent knowledge that is implicitly stored in large databases, data warehouses, the Web or another massive

information repositories. Data mining is a multidisciplinary field that draws on work from statistics, machine learning, pattern recognition, artificial intelligence and more (Han, Kamber, & Pei, 2012).

The two high-level primary goals of data mining are prediction and description. Prediction involves using some variables in the database to predict unknown or future values of other variables; description instead focuses on finding human-interpretable patterns describing the data. Prediction and description can be achieved using a variety of data mining methods:

- Classification: the task of mapping (or classifying) a data item into one of several predefined classes.
- Regression: a function that maps a data item to a real-valued prediction variable.
- Clustering: a descriptive task where one seeks to identify a finite set of categories or clusters to describe the data
- Summarisation: involves methods for finding a compact description for a subset of data
- Dependency modelling: consists of finding a model that describes significant dependencies between variables.
- Change and deviation detection: focuses on discovering the most significant changes in the data from previously measured or normative values.

(Fayyad, Piatetsky-Shapiro, & Smyth, 1996)

Artificial Intelligence and Data Mining

Artificial Intelligence (AI) is usually referred to as “the ability of a machine to learn from experience, adjust to new inputs and perform human-like tasks”. The term AI was first introduced in the 1950s, but with the advancement of Big Data technologies (such as improved computing storage capabilities and increased speed of data processing machines) AI is being revitalised with the power of Big Data. These novel AI systems have improved organisations’ ability to use data to make predictions, as well as reducing the cost of such predictions (Duan, Edwards, & Dwivedi, 2019). Like Big Data, AI is about increasing volumes, velocities and variety of data. When dealing with large volumes of data, AI allows to perform difficult pattern recognition and learning by means of computer-based approaches. AI also contributes to the velocity of data, in that it facilitates quick computer-based decisions. Lastly, AI mitigates variety by capturing, structuring and understanding unstructured data (O’Leary, 2013).

By means of AI it is possible to analyse data trends, provide forecasts, quantify uncertainty, anticipate user’s data needs and suggest courses of actions, Artificial Intelligence is revolutionising decision making with its ability to aid the decision maker in solving complicated and stressful decision problems in real-time, as well as to enable up-to-date information (Phillips-Wren & Jain, 2006). Examples of AI techniques that allow to achieve these goals are rule-based inference, semantic linguistic analysis, Bayesian networks, similarity measures and neural networks (Duan, Edwards, & Dwivedi, 2019).

Machine Learning, on the other hand, is an application of AI based around the idea that machines should be able to access data and use it to learn for themselves without explicitly being programmed (Marr, 2016). A machine learning algorithm is able to support problem solving by learning from a dataset: huge amounts of data are fed into the algorithm, which then uses that data to adjust itself and improve; it then uses what it has learned to solve future problems. Machine learning is fundamentally a way of achieving AI, a way of training an algorithm so it can learn how to accomplish a certain task (McClelland, 2017), such as retrieving insights from data to support organisational decision making.

2.2.2. Big Data lifecycle

In the previous paragraph the concept of Big Data was introduced, as well as the idea that value can be derived by analysing big amounts of data. However, this process is complex and not without its challenges. Several activities are involved in the attempt of turning potentially unstructured data coming from multiple sources into valuable insights for the business.

Lifecycle models provide a structure for considering the operations that need to be undertaken to transform data into knowledge. This structure is often described using the words of ‘Value Chain’, which similarly describes the stages of data processing from the moment it is collected to the moment a decision is made;

additionally, the concept of Value Chain – initially introduced by Micheal Porter in 1980, is aimed at turning the series of activities described into value. Similarly, a data value chain refers to the framework that deals with a series of activities to create value from available data (Bhadani & Jothimani, 2016).

The researcher will use literature from both Lifecycle and Value Chain models, while being aware of the difference between the two. Ultimately, the goal of this paragraph is to provide a description of the series of activities that an organisation would generally undergo to handle data, for which it is not essential to distinguish between the two kinds of models.

The data value chain model of choice for this research is the one identified by Miller & Mork (2012), because they propose a data value chain that aims to manage and coordinate data from data generators to those who consume the information to make decisions. Therefore, it was the emphasis given to the end result of the chain (a decision made) that led to the choice of the model. Some modifications have been made where deemed necessary to reflect the literature reviewed on lifecycle and value chain models.

Miller & Mork (2012) distinguish between three major phases, namely *Data Discovery*, *Data Integration* and *Data Exploitation*. The value chain is illustrated in Figure x and the phases are described below:

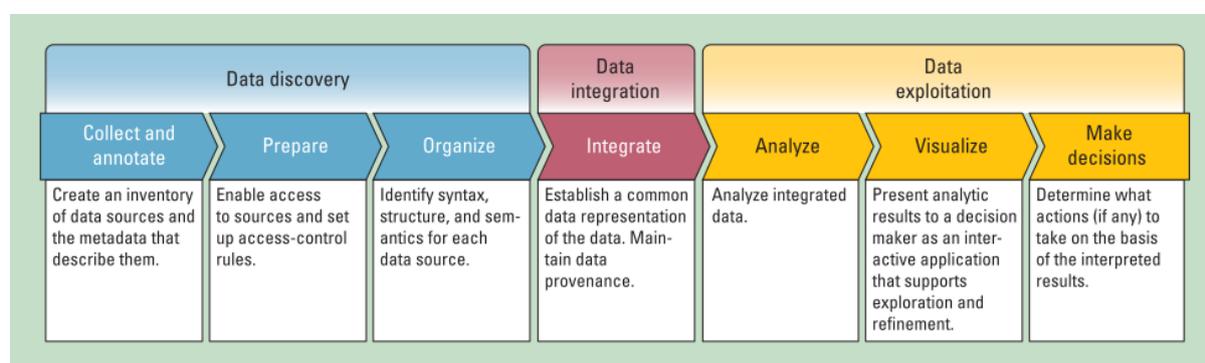


FIGURE 4 – DATA VALUE CHAIN (MILLER & MORK, 2012)

1) Data Discovery

Before performing analysis with data and use the information derived from it to support informed decision-making, an organisation needs to gather the necessary data resources. The data discovery phase involves the collection and inventorying of the data assets, but also its preparation to the analysis. The Data Discovery phase is comprised of the following sub-phases:

1a) Collection and Annotation: the first step is to collect raw data from all possible – and relevant – data sources (Bhadani & Jothimani, 2016). This phase is also focused on turning the collected unstructured data into structured data, by associating valid metadata to the unstructured data: thus, an inventory of the data sources and the metadata is created, which also pays attention to the quality of the sources (in terms of completeness, validity, consistency, timeliness and accuracy).

1b) Preparation: the data sources are copied into a shared system and access control rules are set up to impose restriction on the data use and guarantee the security and privacy of the data. The data centre where the data is transferred to helps in collecting, organising and managing data (Bhadani & Jothimani, 2016).

1c) Organisation: organisational choices are made about the data's syntax, structure and semantics by the data source developer; this information is then made available through a schemata or a metadata repository.

2) Data Integration

The integration phase serves to combine the available data, which came from heterogeneous sources, into a uniform, common representation. This facilitates users accessing and querying such data as if they were accessing only one data source (El Arass & Souissi, *Data Lifecycle: From Big Data to Smart Data*, 2018).

Combining different sources not only serves to create mappings between data sources, it also delivers new, undiscovered information.

Data Integration, according to Bhadani & Jothimani (2016), is one possible way that data can be pre-processed to get rid of redundancy, noise and inconsistency in the data. They therefore distinguish between the sub-phases of *Integration*, *Cleaning* and *Elimination of Redundant Data* within the major phase of *Data pre-processing*. El Arass & Souissi (2018) also include in their review of Data Lifecycle models separate *Filtering* and *Enrichment* phases in which, respectively, data of poor quality is filtered out of the process and additional information is added to enrich the data currently being used in the cycle. The model of Miller & Mork (2012) does not make such distinction explicit. The researcher chose to modify the model, noting that at this stage of the value chain some cleansing of the data is occurring to increase the quality of data. An additional step of *Cleansing* was added under the *Data Integration* phase. In the way the phases were ordered in the model, the *Cleansing* step is executed after the *Integration* of data. This was done consistently with the review of El Arass & Souissi (2018), but arguably the two phases are not necessarily executed in that order; they might be complimentary as discussed by Bhadani & Jothimani (2016).

3) Data Exploitation

When data gets to the Data Exploitation phase, it has been gathered and integrated and it is now ready to be analysed and visualised in order to convey insights to decision makers – who can use the generated information as a basis for – informed – decision-making. These three steps are described in more detail below:

3a) Analyse: in this phase, the raw data is analysed to draw information and knowledge from it (El Arass & Souissi, *Data Lifecycle: From Big Data to Smart Data*, 2018). This phase also includes maintaining metadata and the provenance between inputs and results so that another analyst can recreate the same results and strengthen their validity. Furthermore, one of the most important steps of data analysis is selection of appropriate techniques for data analysis (Bhadani & Jothimani, 2016). A few examples of Big Data analytics algorithms have been described in paragraph 2.2.1.

3b) Visualise: the analytic results are displayed in a clever and intelligent way – in the form of a static report or interactive application, and presented to decision makers; the goal is to turn meaningful information in a format that decision makers can easily understand and consume to make decisions (El Arass & Souissi, *Data Lifecycle: From Big Data to Smart Data*, 2018).

3c) Make decisions: during this stage it is determined what action is necessary given the visualised results. The details of a particular problem have been analysed and visualised so that informed decisions can be made (Bhadani & Jothimani, 2016). Supporting documentation describe how analysts obtained the results and should include provenance information to the original sources, quality annotations, integration mappings and analysis metadata.

El Arass & Souissi (2018) argue that most lifecycle models include a *Destruction*, and *Achieving* phases after the data analysis is complete and the information has been generated for decision makers to use it. An organisation has a choice between: deleting the data when it is successfully used and will become useless without added value; store the data long-term for possible future usage. Possibly, Miller & Mork (2012) decided to end the value chain with the decision making phase because arguably that is the stage where value is derived from data. There is no added value from disposing of the data, whereas the value that comes from archiving the data is uncertain.

Nonetheless, based on the analysis of El Arass & Souissi (2018) the researcher decided to add the additional phases of *Data Storage* and *Data Disposal* were added to reflect the actions of archiving and disposing of the data after the exploitation of data.

The resulting value chain model is depicted in the figure below:



FIGURE 5 – BIG DATA LIFECYCLE

Whether the data is big or small, at each step there is work to be done and phase-specific challenges to be addressed with Big Data (Labrinidis & Jagadish, 2012). These challenges, specifically the ethical ones, will be discussed in the next chapter.

2.4. Big Data Ethics

Big Data provides us with huge opportunities for understanding and predicting customer behaviour, for managing supply chains, for creating new products and services or enhancing existing ones: it brings the potential to improve our private and public lives overall. Unfortunately, these opportunities are coupled to significant ethical challenges and potential risks for firms and their customers (Herschel & Miori, 2017). The use of more data, often personal and potentially sensitive, and the reliance on algorithms to analyse the data to shape choices and make decisions, pose ethical issues of, to mention a few, fairness, responsibility and respect of human rights (Floridi & Taddeo, 2016). Gartner asserts that combining Big Data with sophisticated analytics capabilities increases the risks of business ethics violations (Herschel & Miori, 2017).

Big Data analytics raises an ethical debate that interests the whole chain of data activities that result in informed decision-making. Starting from the moment data is collected, the quality of data gathered can be affected by the fact that it comes from multiple data sources in different contexts, as well as the fact that it often unstructured data coming from social media sites. When this data is used to identify patterns from groups, the derived insights might have errors and biases for individuals who do not conform to group characteristics, leading to discriminatory situations. Also, when poor quality data is used as an input to complex and hard to understand algorithms, problematic situations may arise if that incorrect data and/or algorithm led to unethical decisions. Furthermore, when Big Data informs decision-making, it is hard to justify how decisions are made, hence questions about responsibility arise (Asadi Someh, Breidbach, & Davern, Ethical Implications of Big Data Analytics, 2016).

The new insights and predictions that come from analysing Big Data are already starting to have an impact on citizens, governments and companies. Big Data is being adopted in an increasing number of fields and activities ranging from dating to hiring, voting and identifying terrorists. It has been happening so quickly that most people are not aware of both the scale and speed of these transformations (Richards & King, 2014). In this rapidly changing society, the study of ethics is particularly important. While new technologies, Big Data among them, have brought us many benefits, there is the risk that some people or organisations might exploit them for personal gain; these behaviours raise ethical concerns that need to be addressed (Quinn, 2004).

Data ethics was born to respond to these challenges, with the goal of maximising the value of data science towards society. This field can build on the foundation provided by business ethics – which addresses ethical challenges in the business environment, together with the knowledge developed by computer and information ethics – which focuses on the challenges posed by digital technologies (Floridi & Taddeo, 2016).

Defining Big Data Ethics

Floridi & Taddeo (2016) provide the following definition of Data ethics:

“Data ethics is the branch of ethics that studies and evaluates moral problems related to data (including generation, recording, curation, processing, dissemination, sharing and use), algorithms (including artificial intelligence, artificial agents, machine learning and robots) and corresponding practices (including responsible innovation, programming, hacking and professional codes), in order to formulate and support morally good solutions (e.g. right conducts or right values).”

They explain that the ethics of data focuses on ethical problems concerning:

- 1) Data (in a strict sense): this refers to activities such as the collection and analysis of large datasets. Key issues concern: re-identification of individuals through data mining, linking and merging large datasets; identification of types of individuals, which may lead to group privacy violations when people are discriminated and targeted based on the group that they belong to; lack of transparency of the data analytics process.
- 2) Algorithms: this refers to the increased complexity and autonomy of algorithms (especially in the case of Artificial Intelligence and Machine Learning). The ethics of algorithms raise challenges of moral responsibility and accountability of the data 'actors' involved in the process of turning data into information used for decision making purposes.
- 3) Practices: this refers to the responsibilities of people and organisations in charge of data processes, strategies and policies. The goal is to define ethical practices that ensure a responsible use of new technologies and protect the rights of individuals and groups.

The authors also argue that data ethics must investigate the whole space of data ethics throughout the three above-described axes; the diverse set of ethical implications of data science should therefore be addressed by a consistent and inclusive framework.

(Floridi & Taddeo, 2016)

One of the activities of business ethics involves the study of business practices, which are investigated from a moral point of view to determine whether the activities under the eye of the investigation are moral or immoral. For example, a study of hiring or firing practices could be part of business ethics if the intent was to determine whether or not discrimination could be detected, or whether a practice designed to end discrimination was working successfully (De George, 1987). Similarly, the aim of Data ethics is to study business practices involving data. For example, given a specific series of activities that a business executes with data – which resembles a data lifecycle, data ethics should work to detect potential ethical violations within those activities. A company may have rules that apply to business practices that ensure that employees abide by laws (Farlex, n.d.); a similar line of reasoning may be applied to data ethics. Rules should be introduced in the business to ensure the compliance to data laws and regulations.

The researcher sees a direct link between the fields of business ethics and (big) data ethics and since existing literature does not provide a definition of Big Data ethics specifically, a definition of Big Data ethics is proposed based on the definition of data ethics provided above, as well as the previously defined concepts of Ethics and Business Ethics and the way these fields are interrelated:

“Big Data ethics is a branch of business ethics that studies ethical problems that arise in the business environment when using Big Data and algorithms for data analysis. Its goal is to develop moral rules, standards, or practices that support moral decision-making based on Big Data analytics”.

Since this research focuses on the uses of Big Data specifically within the context of commercial businesses, the definition of Big Data ethics is being placed in the bigger field of business ethics. The mentioned moral rules, standards and practices are said to support a moral use of Big Data analytics: this refers to the whole process of handling data from its collection to the moment a decision is made based on the information produced by the analysis itself. Furthermore, the definition discusses Big Data analytics in terms of its ultimate goal, which is to aid decision makers in making informed choices, however it also underlies all the activities described in value chain and lifecycle models of data.

In this section we have analysed how Big Data has a dual facet: on one hand, it is seen as a powerful tool to solve societal issues by offering insights into areas as diverse as cancer research, terrorism and climate change. On the other, Big Data enables invasions of privacy and violation of basic human rights and thus raises difficult ethical questions (Boyd & Crawford, 2012).

It is therefore necessary to analyse what practices and activities enable the violation of ethical principles and ask which systems and laws are in place to regulate them.

2.3.1. Applying ethical values to Big Data

While business innovators are excited about the potential benefits they can create from Big Data and its related technologies, the size, variety and velocity of these newly generated information raises questions about the implications of acquiring, storing and using large quantities of personal data – which concerns people’s characteristics, behaviours, preferences and location, among other things.

These questions that are being raised are ethical, in that they relate to the ethical values deeply established in ourselves, and how such values are applied in the creation of knowledge, products and services. These values can inform us on how we should execute the design of algorithms and operations on these massive amounts of data available; they also allow us to weigh the benefits of Big Data against the risks of unintended consequences.

(Davis & Patterson, 2012)

Before ensuring that Big Data is used to bring benefits to society and minimise risks, first it is important to agree on high-level principles that can “help condense complex ethical issues into a few central elements which can be clearly understood and agreed upon by people from diverse fields and sectors” (Whittlestone, Nyrup, Alexandrova, & Cave, 2019). However, while these principles occupy a central role in applying ethics, they are not enough to ensure that society can reap the benefits and mitigate the risks of new technological phenomena such as Big Data. Beauchamp and Childress (2001) suggest that principles should be taken as guidelines, which need to be made specific for use in policy and decision-making. They also elaborate that in order to be action-guiding, principles need to be accompanied by an account of how they apply in specific situations.

The operation of applying ethical principles to specific situation is not an easy one, in that there is the risk of encountering ambiguous situations. To give an example of this, consider the efforts of organisations to use data analysis to recommend their customers a book to read, a movie or TV series they might be interested in, or a new shirt to buy based on previous purchases. It is hard to identify any ethical violations in this use of data. However, let us consider the same analysis algorithms that, in 2012, predicted that a man’s daughter shopping at Target was pregnant – even before the man and his family knew she was, based on the fact that she had purchased unscented wipes and magnesium supplements; the company was sending the man coupons for baby clothes and maternity wear for his teenage daughter to wear. The question arises of whether it is ethical to send such coupons based solely on the data that had been analysed by the algorithm. In these two examples, the difference is in the context: when predictive analytics is executed on sensitive categories of data such as race, sexuality and health, ethical dilemmas are more likely to arise.

In section 2.1.1 the five high principles of Autonomy, Beneficence, Nonmaleficence, Justice and Privacy and data protection have been discussed, as introduced by Beauchamp & Childress (2001) and Wright (2011). These principles can be applied to the context of Big Data to result in specific situations of ethical dilemma that concern Big Data. In this process, the sub-values embodied in the major principles are related to ethical aspects found in literature concerning Big Data, Data Mining and Artificial Intelligence/Machine Learning.

1) Respect for autonomy

1a) Autonomy:

Algorithms can appear ethically neutral because they affect how we see the world and modify its social and political organisation. Algorithmic activities, like profiling, reconceptualise the world in new ways and motivate actions based on the insights it generates. Algorithms can nudge the behaviour of data subjects and human decision-makers by filtering information. For example, through the use of personalisation algorithms groups of people within a population are offered different content, information, prices etc. according to a particular attribute (such as their preferences or ability to pay). This phenomenon draws a thin line between supporting and controlling the decisions of users. In filtering the information presented to the user based on the algorithm’s understanding of their preferences and behaviours, the subject’s autonomy is disrespected, because their choices might not be deliberated on the basis of their desires and plans, but rather on the basis of the interests of a third party (namely, the organisation that makes use of personalisation algorithms to

target specific groups of customers).

We thus observe a paradox: personalisation should facilitate the decision-making process of the subject by showing them only what is relevant to them; however, by not providing the subject the complete picture of information, they can be pushed to make the action preferred by the third party rather than guided by their own preference. The user's actions are somewhat manipulated and used by the third party to achieve their own goals – which might take, for example, the form of revenue. These organisations act on the basis of their plans and desires (e.g. targeting an individual with specific advertisements to potentially earn money from that individual), and the users find themselves in a position where, once the data has been collected, they lose control of their decision making process.

(Mittelstadt, Allo, Taddeo, Wachter, & Floridi, 2016)

1b) Dignity:

Company's marketing efforts of personalisation could violate human dignity if users are targeted and treated differently based on their age, gender, race and economic situation. Companies that possess great amounts of user data might know the preferences, economic possibilities and personal information of individuals transacting inside and out. They may, for example, use the information at their disposal to usurp the entire value surplus available in the transaction by pricing goods or services as close as possible to the individual's reservation price – meaning the highest price the user is willing to pay (Tene & Polonetsky, 2013). Such behaviour does not respect human dignity because it implies that those with a better economic situation are being valued more than those with a worse economic situation: those with better financial possibilities will be shown a higher price for a product than those that can afford less.

1c) Informed consent:

Thanks to the GDPR regulation, consent has to be provided by users to companies who wish to collect and retain personal data. If an individual has unambiguously given his or her consent, data processing is legitimate. As long as this process involves personal data in a strict sense the law applies without reservation. However, once the data has been anonymised, the GDPR does not apply anymore and the individual loses control over the data process.

Furthermore, these guidelines demand that the reason for collecting personal data should be made clear to the individual prior to collecting it. However, problems arise because it is likely that not even those who operate with data know what the data will exactly be used for (Wahlstrom, Roddick, Sarre, Estivill-Castro, & deVries, 2006). Furthermore, Big Data makes use of passive technologies, such as location-based information from mobile phones or data from sensors. Even if the individual has initially given permission to gather data, on the long term they may no longer be aware that data is currently being collected about them, because these services do not ask for permission every time contextual data is gathered (Nunan & Di Domenico, 2013). Thus, the users are giving their consent for personal data usage, potentially without full understanding what they are agreeing to (Wahlstrom, Roddick, Sarre, Estivill-Castro, & deVries, 2006). Even if the individual provides information that he is comfortable sharing, the interferences drawn from the data can reveal information that he does not want revealed and which may be harmful to him. Most individuals are not fully aware of the possibility of revealing this more sensitive information from their personal data and some, if they knew, would view it as an unauthorised appropriation (Cary, Wen, & Mahatanankoon, 2003).

1d) Social solidarity, inclusion and exclusion:

Despite the wide spread of the internet and mobile devices among the population, cost and knowledge reasons can impede certain groups of people from accessing them. This leads to a representativeness problem, meaning that the population represented in the data is only a small group of people. This problem is exacerbated by the fact that these companies might be led to believe that the large size and volume of Big Data that they collected is representative and not random. However, the quantity of data does not guarantee its quality. Studies have shown that the social media population is far from being representative of the entire population, or even the Internet user population, and that the representativeness problem is both in the age structure and in regional division. For example, a survey conducted by CINIC in 2014 showed that nearly 70% social media users are under 30 years old; thus, the social media population is a small sample mainly

populated by young people. Another study conducted in 2012 interviewed 1802 American Internet users and showed that only 16% of Internet users have a Twitter account, and the majority of them are African-American, urban residents, and young people between 18 and 29 years old (Liu, Li, Li, & Wu, 2015). Ethical questions arise when only those groups that have an internet connection, a mobile phone and/or social media accounts, are offered certain products and services, while those that are excluded from the information society might be missing out on offers or opportunities.

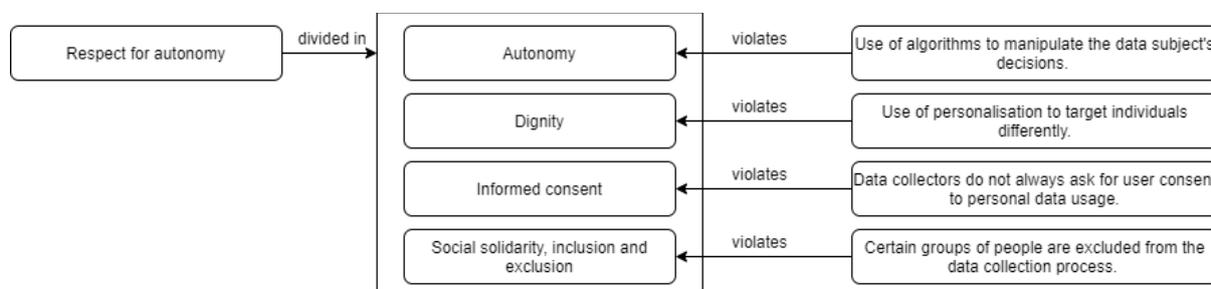


FIGURE 6 – UNETHICAL ACTIVITIES FOR THE RESPECT FOR AUTONOMY PRINCIPLE

2) Nonmaleficence

2a) Safety:

Predictions and knowledge inferred from the analysis of Big Data may cause harm to consumers. Zwitter (2016) states that the phenomenon of data scientists letting algorithms search for correlation themselves increases the danger that the algorithm might establish random correlations based on co-occurrences, which might favour certain people and penalise others (Zwitter, 2014). Furthermore, if the representations of analysis results are unclear, the human may interpret them his own way, possibly introducing bias in the results (Fayyad, Piatetsky-Shapiro, & Smyth, 1996). Depending on the specific application that the information is used for, this can be more or less harmful for the user and may lead to ethical issues: for example, if an individual's insurance application is refused on the basis of some personal characteristic of the user, the person is put into a disadvantageous position and arguably the Big Data system isn't protecting his or her economic interests. Another example comes from O'Neil's book "Weapons of Math Destruction", in which the author describes the race of American universities to outrank each other, chasing measures imposed by the mathematical models that determine the ranking themselves – measures such as acceptance rates, the percentage of alumni who donate to the school and the students' SATs scores. This goes to the detriment of students who face increasing college costs without any improvements to the actual quality of the education received, and that see their right to education being threatened (O'Neil, 2016).

2b) Isolation and substitution of human contact:

This problem is not discussed in the Big Data ethics literature because it is not strictly related to Big Data activities, but rather to those new information technologies – such as new communication tools – that might create social isolation by substituting face-to-face contact for virtual communications. It will be therefore excluded in the further analysis stage of this research.

2c) Discrimination and social sorting:

Big Data supporting technologies (such as Data mining and Artificial Intelligence), are tools with a discriminatory nature: they allow social sorting and segmentation which could have unfair effects on the population (Wright, 2011). In fact, Big Data analytics places individuals into pre-determined categories; society is compartmentalised into groups and the individuals susceptible to disease, crime, or other socially stigmatising characteristics or behaviour might pay the consequences. Wright (2011) states that these technologies allow social sorting and segmentation which could have unfair effects on the population. Surely predictive analytics can be used to benefit society, but when it is executed on sensitive data that regards the health, race or sexuality of individuals, it might help perpetuate old prejudices: the wealthy and well-educated will get the fast track, the poor and underprivileged will face more adversities than before. Predictive analytics

risks becoming a self-fulfilling prophecy that accentuates social inequalities (Tene & Polonetsky, 2013). When data is analysed to discover patterns, the ladder can be used to build profiles of characteristics of behaviour of individuals. This practice, called *profiling*, identifies correlations and makes predictions at a group-level (Mittelstadt, Allo, Taddeo, Wachter, & Floridi, 2016): discrimination issues may then arise if individuals are judged on the basis of the attributes of the group to which they belong, rather than on the basis of their own particular characteristics (Wahlstrom, Roddick, Sarre, Estivill-Castro, & deVries, 2006). The classification of individuals into groups based on race, ethnic group, race, gender and social and economic status could result in the offering or restriction of special treatments or services to individuals or groups (Asadi Someh, Breidbach, & Davern, Ethical Implications of Big Data Analytics, 2016).

It is clear why there is an ethical interest in a correct and accurate use of data: the consequence is potential discrimination, which might not directly affect individuals, but can have an impact on local communities in terms of social stigma and inadequate provision of services (Mantelero, 2017).

All users should receive equal advantage from the use of Big Data, thus it is important for companies that make use of Big Data systems to look out for potential discrimination in the data, as well as in the algorithms, in order to reduce the risks of overlooking particular groups of people (Stoyanovich, Abiteboul, & Miklau, 2016).

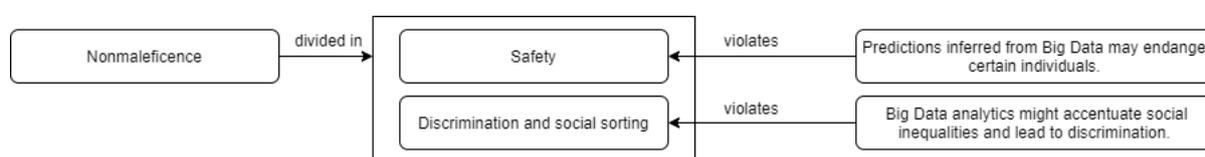


FIGURE 7 – UNETHICAL ACTIVITIES FOR THE NONMALEFICENCE PRINCIPLE

3) Beneficence

3a) Imbalance of power:

Data collection and analysis efforts tend to only be advantageous for the organisations that embark on such exercises. While from the outside companies might argue that their doing favours a broader community, this isn't necessarily true. Decision makers use the outcomes generated by Big Data analytics to take decisions that affect individuals and groups without allowing them to participate in the process, which causes a situation of imbalance between the data gatherers and the data subjects: individuals may be giving away their personal data, and not directly see value returned back to them. Furthermore, the fact that individuals do not know the complexity of the Big Data process means that they are not aware of the potential prejudices that underlay its use, and thus they are not in the position to object to the discriminatory use of personal information by data gatherers (Mantelero, 2017).

Chessel (2014) argues that the results of data analysis should be equitable to all parties, and everyone should be fairly compensated. However, this imbalance of power prevents commercial organisation from distributing value to the users and actively doing good to them.

The problem of imbalance of power was added by the researcher under the category of Beneficence, to refer to a situation in which one of two parties holds the power in the relationship: this value makes explicit one of the constraints that companies are subject to and prevent them from living by this ethical principle.

3b) Universal service

This value is not strictly related to Big Data systems because it refers to the obligation imposed on operators of electronic communication networks to provide a minimum set of services to all citizens. Ultimately, the amount of people that have access to these services determines the population that is included in the information society, and thus has an impact on Big Data systems and the representativeness of data collected. However, since it is not within the obligations of commercial companies collecting and using data for analytics purposes to provide these services, this value will be discarded in the further analysis of the research.

3c) Accessibility

The issue of accessibility is related to the value of social solidarity that has been discussed before. As in the

case of social solidarity, when it comes to Big Data the lack of accessibility of devices or of an internet connection can prevent companies from collecting data about those who are incapable of accessing websites or mobile devices. Wright (2011) argues that the market tends to overlook the needs of the disabled and of senior citizens, which mostly affects the accessibility of websites, digital television, phones, emergency services and public information terminals. This can cause data to not be representative of the whole population which leads to a lack of fairness of the sampling of the population and potential inaccuracy of the results of the data analysis. With 15% of the EU population suffering from some form of disability, they represent a mass market that is being excluded from Big Data systems (Wright, 2011). This value, while being related to Big Data, is not something that is in control of the companies that collect information about individuals – despite this being not representative of the entire population, but rather something they should be aware of. This value is therefore not considered in the further analysis of the research.

3d) Value-sensitive design

Big Data analytics is a human-supported process and it is therefore subject to the error and bias of to human that might affect the data, measures, the design of the algorithm and the analysis. Potential errors could start occurring in the collection and storing of data in the database due to human intervention in such phases (Asadi Someh, Breidbach, & Davern, Ethical Implications of Big Data Analytics, 2016). Furthermore, often the automation of decision-making through the use of Big Data analytics is justified by an alleged lack of bias in algorithms. However, algorithms are designed by humans, and so they reflect the values of its designer. Development is not a linear path: there is no objectively correct choice to make at any stage of development; as a result, “the values of the author [of an algorithm], wittingly or not, are frozen into the code, effectively institutionalising those values”. The visualised outputs of algorithms also require interpretation, and the human’s “unconscious motivations, particular emotions, deliberate choices, socio-economic determinations, geographic or demographic influences” might influence the way correlations are interpreted (Mittelstadt, Allo, Taddeo, Wachter, & Floridi, 2016).

Thus, Big Data and its supporting technologies are arguably not ethically neutral. The human design can affect the whole lifecycle of data, from moment data is collected (e.g. certain attributes might be deemed not relevant from the human, thus causing them to be excluded from the collection and/or the analysis), to the moment data is analysed (e.g. the algorithm might reflect the bias of the human who designed it), to the data usage phase (e.g. results might be misinterpreted or misunderstood by the human).

3e) Sustainability

This value cannot be directly related to Big Data and Big Data analytics because the word sustainability is used as an attribute to a project to refer to the condition where it can be sustained into the future: this refers both to financial support for the project, as well as the consideration for the environment. Wright (2011) states that developers should be aware of the consequences of using certain materials for the production of new technological products, be aware of the problems of depletion of natural resources and opt for more recyclable materials when possible. Because sustainability is not an ethical problem faced by users of Big Data and related supported technologies, it is excluded from now on from the analysis of Big Data ethics.

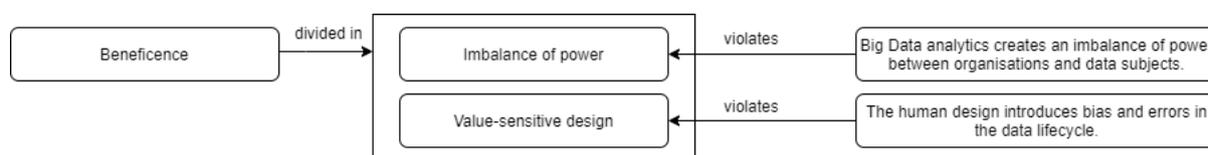


FIGURE 8 – UNETHICAL ACTIVITIES FOR THE BENEFICENCE PRINCIPLE

4) Justice

4a) Equality and fairness

The principle of justice refers to treating other humans fairly and equally. Problems of lack of fairness and equality in Big Data systems have already been discussed in the *Discrimination and social sorting* section, as

well as in *the Social solidarity, inclusion and exclusion* section of this sub-chapter. Despite the overlap with these values, the way the principle of justice relates to the Big Data context will be discussed separately.

Stoyanovich, Abiteboul & Miklai (2016) discuss problems of fairness in using Big Data analytics for classification of users. They distinguish between individual fairness and group fairness. Individual fairness states that two individuals who are similar with respect to a particular classification task should be classified similarly, while group fairness states that the proportion of members of a protected group who are classified positively should be statistically indistinguishable from the proportion of members of the overall population. The authors fear that if these goals are not pursued by Big Data technology, inequalities will increase.

The common practice of *personalisation*, which consists in segmenting the population so that different services can be offered on the basis of behaviours, habits and personal characteristics, raise questions of fairness and equitability: in fact, this practice may segment the population into groups so that only some of them are worthy of receiving some opportunities or information – thus reinforcing existing social (dis)advantages between groups of people.

Violation of fairness and equality also arise from the lack of representativeness of data (unfair sampling) as well as bias introduced in the design of the algorithm by the human and from the misinterpretation of analysis results. If decisions are made that do not take these errors and bias into considerations, harm or disadvantage might be caused to certain individuals or groups of individuals. For example, O’Neil describes in her book “Weapons of Math Destruction” that biased algorithms cause insurance companies to charge each individual for the highest price they will tolerate, thus treating individuals differently based on their economic possibilities. In some cases, the use of biased systems will cause some people to even be denied that insurance, or a job or a loan on the basis of their credit scores. These examples make evident how Big Data systems might violate the principle of justice and might exacerbate the inequalities between groups of people.

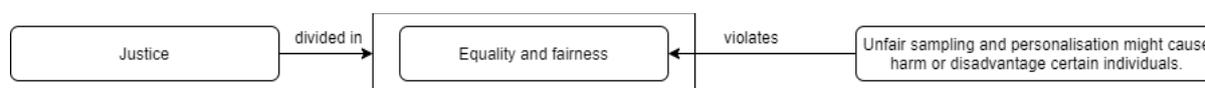


FIGURE 9 – UNETHICAL ACTIVITIES FOR THE JUSTICE PRINCIPLE

5) Privacy and data protection

5a) Collection limitation and retention

One of the fundamental statements of the GDPR requires organisations to limit the collection of personal data to the absolutely necessary for carrying out the purpose for which the data is collected in the first place. It also requires organisations to delete data that is no longer being used for the purposes for which it was collected. Moreover, it promotes a more considerate behaviour from organisations when it comes to retaining individuals’ personal data. Tene & Polonetsky (2013) point out that organisations today collect and retain personal data through multiple channels including the Internet, mobile devices, sensors, and emails; data is collected either directly from individuals or through third parties (from semi-public sources like Facebook to public sources like the government). With this in mind, the authors argue that data minimisation is not the market norm. Despite the GDPR having been introduced to address this problem, it is still uncertain how organisations apply the data minimisation principle in practice. Difficulties arise, for example, when executing Data Mining on a set of data. Wahlstrom, Roddick, Sarre, Estivill-Castro & deVries (2006) explain that it is impossible to accurately define the purpose of a data mining exercise, as it is intrinsically related to the information it discovers. Also, oftentimes data is collected without a precise objective in mind (Fayyad, Piatetsky-Shapiro, & Smyth, 1996), causing unnecessary data to be collected and thus potentially violating the GDPR principle. As a last point, data mining is conventionally executed over large amounts of historical data. Thus, while on one hand organisations are encouraged to reduce the retaining periods of data, they might not want to do so in order to not undermine the outcomes of their data mining efforts.

5b) Data quality

When dealing with data of the volume of Big Data, it is difficult to manage the quality of data, especially

considering that data is sourced from multiple sources in different contexts. Collecting vast amounts of data from diverse – possibly external, sources may cause problems because the quality of the data cannot be assured and may be noisy, obsolete, inaccurate or incomplete (Wahlstrom, Roddick, Sarre, Estivill-Castro, & deVries, 2006).

Especially the analysis of data coming from social media sites is unstructured (the quality of which is generally low), might result in the discovery of patterns affected by considerable errors and biases (Asadi Someh, Breidbach, & Davern, Ethical Implications of Big Data Analytics, 2016): the human operating with the data might make assumptions to fill in the missing values, leading to a biased dataset (Fayyad, Piatetsky-Shapiro, & Smyth, 1996) and, possibly, to bad decision making.

When interferences are made about particular individuals, based on data that may be of poor quality, unethically discriminatory decisions might result (Asadi Someh, Breidbach, & Davern, Ethical Implications of Big Data Analytics, 2016), with repercussions for the data subject (Wahlstrom, Roddick, Sarre, Estivill-Castro, & deVries, 2006) and especially for those individuals who do not conform to group characteristics. Thus, it is important that data is collected in a timely manner and kept updated, to prevent skewed results to be produced and consequently wrong decisions to be made.

5c) Purpose specification

The problem of purpose specification is strictly related to informed consent: in fact, in order to be meaningful consent must be specific to the purpose (or context). And yet, while it is within the GDPR guidelines to always specify the purpose of the data collection during the data collection phase itself, by its very nature the analysis of Big Data brings surprising correlations and produces results that resist prediction (Tene & Polonetsky, 2013). In fact, companies frequently use data analysis techniques on historical data – with the purpose of acquiring new knowledge about individuals and groups, meaning that the data collected for one purpose is likely being used for another purpose. This makes it nearly impossible to allow the customer to (1) have the right of giving informed consent for each use of his data (as previously discussed in the Informed Consent section) and (2) know how the data being collected will eventually be used (Cary, Wen, & Mahatanankoon, 2003). The problem is intrinsic to the technology, which doesn't allow the purpose of the analysis to be known until it has successfully revealed some previously unknown information. In other words, the purpose of activities such as data mining is strictly related to the information they discover (Wahlstrom, Roddick, Sarre, Estivill-Castro, & deVries, 2006).

Because the personal data that individuals may have willingly granted for use in one context is often subsequently mined for purposes other than the one for which data was granted, we can question whether the data subjects are being treated fairly (Tavani, 2004) and whether there are some ethical violations happening – besides violations of the EU data regulation.

5d) Use limitation

While Wright (2011) discusses the issues of Collection limitation and Use limitation separately, there is an overlap of meaning between the two. Furthermore, the issue of limiting the collection (and consequently the use) of data has been already discussed in a previous section. Thus, the researcher has chosen to discuss the limitation of use in terms of the sharing of data from original data collectors with other parties.

Ethical questions arise when data is shared or sold between organisations, because these operations make it harder to determine how the data was collected, what new usage will be made out of it and whether the person who provided initial consent to the usage of their personal data agrees to the usage that the new party will make of it. These organisations often assume that the users consent applies to any future use of their data and to not make an effort to inform them of the movements of the data or allow them to opt-out of the practice. Aside from the information collected publicly from sources such as the Internet, much information is bought from private sources and can include credit history, financial information, employment history and possibly some medical information (Cary, Wen, & Mahatanankoon, 2003). The sharing and buying/selling of data ultimately makes it hard to enforce a limit on the usage of data and creates ambiguous situations in which it is hard to tell whether any regulation is being violated.

5e) Confidentiality, security and protection of data

As the volume of data increases, new opportunities for data breaches are created, the lack of protection of data would mean, for the organisation that possesses it, violating ethics (White & Ariyachandra, 2016).

The value of security is not dealt with in this research because authors such as Alshboul, Wang & Nepali (2015) as well as Ye, Cheng, Yuan, Xu, Gao & Cheng (2016) have already discussed in depth security in the context of Big Data, in terms of threats and challenges throughout the Big Data lifecycle – including risks of loss of data, unauthorised access or disclosure of data, among others; in order to avoid repetitions and keep the focus of the research towards providing new knowledge to the existing body of literature, the aspect of Security in Big Data will be excluded from further analysis.

5f) Transparency

For Stoyanovich, Abiteboul & Miklau (2016) transparency means being able to verify and audits datasets and algorithms for fairness, robustness, diversity and non-discrimination. Transparency is generally desired because algorithms that are poorly predictable or explainable are difficult to control, monitor and correct (Mittelstadt, Allo, Taddeo, Wachter, & Floridi, 2016). Transparent algorithms allow to render complex decision-making processes both accessible and comprehensible (Mittelstadt, Allo, Taddeo, Wachter, & Floridi, 2016). A lack of transparency, on the other hand, implies that the human does not know how the Big Data system makes decisions, which ultimately makes it difficult for them to identify the fairness to individuals involved, robustness and non-discriminating ability of the system itself. Transparency implies that the humans that operate with the data analysis system need to understand how it will act in different circumstances; they know how the system works and are able to prevent it to behave in an unexpected manner (World Economic Forum, 2019). However, when a data analysis system, such as an Artificial Intelligence algorithm, is too complex, it might be difficult even for the engineers who designed it to decipher why the machine made a certain decision (Business Ethics Briefing, 2018). Furthermore, information about the functionality of algorithms is often poorly accessible: proprietary algorithms are intentionally kept secret for the sake of competitive advantage. The challenge of transparency is therefore tied to the opacity of algorithms, which refers to the fact that “if one is a recipient of the output of the algorithm, rarely does one have any concrete sense of how or why a particular classification has been arrived at from inputs” (Mittelstadt, Allo, Taddeo, Wachter, & Floridi, 2016).

Here thus emerge ethical concerns: not only the trust in the system decreases, but it is also hard to determine who is responsible in the event of any damage to individuals (Business Ethics Briefing, 2018). As data analysis technologies become more opaque, human interaction with the system becomes more complicated, and the decisions made consequently lack transparency; this phenomenon leads to the challenge of identifying ethical violations happening within the Big Data system itself.

5g) Individual participation and access to data

When it comes to individual participation to data, what often happens is that users are unaware of the forms of data analysis – which allow to infer predictive information about groups or people, as well as the impact that the information collected or generated may have on themselves and/or another group. Furthermore, decision makers use the outcomes generated by Big Data analytics to take decisions that affect individuals and groups, without allowing them any participation in the process (Mantelero, 2017). Tene & Polonetsky (2013) agree with this statement and say that even when organisations comply with the law and grant the right of access to data to individuals, they implement the data protection directives narrowly: in fact, they provide individuals with little useful information and are seldom willing to share the wealth created by this personal data with the involved individuals. Also, if organisations were to fail to properly track the sources of data, they would be unable to provide information about individuals’ data or, when requested, to delete it altogether. The individual access to data is also made difficult by the fact that, despite individuals saying that they would like to exercise control over their data, in practice they do not take advantage of the control that they already have through regulations such as the GDPR. For example, individuals may rarely read terms and conditions carefully before consenting to them, or they rarely think about the ways that data controllers may gather information about them. Their interest in control is only spiked when something goes wrong (British Royal Academy, Royal Society, techUK, 2018).

As for the right of an individual to have their personal data erased, this can prove to be a challenge in the case

of Big Data. Nunan & Di Domenico (2013) state that Big Data has the ability to rewind and fast-forward people's lives, but in doing so it may remove the ability for individuals to be forgotten. In practice, the activity of transferring data between companies makes it difficult to keep track of data movements; thus, it might prove difficult to permanently erase user data from these companies' databases, because they might request the deletion to a party that in that moment in time might not possess such data anymore – because it has been sold or shared with a third party.

5h) Anonymity

Big Data can be used to make decisions about a population, solely based on quantitative information. When personal aspects of data are removed, however, individuals or groups are compartmentalised: ignoring certain personal characteristics from the analysis might lead to different decisions; on the other hand, however, keeping too many aspects of data regarding the individual works against the attempt to anonymise data and thus the attempt to preserve the user's privacy (White & Ariyachandra, 2016).

To complicate this situation even further, there is the fact that sometimes companies' efforts to anonymise user data to guarantee their privacy are simply not enough. One example of this is evident in Web search engines: users' queries may expose facets of their life, interests, personality, sexual preferences and health issues they might not want to share with everybody. And yet, two New York journalists were able to spot several queries, originating from the same user and referring to the same last name or specific locations, that could be linked to a senior woman – who then confirmed to have issues these queries. Thus, trying to anonymise data by replace a user name with a number is not enough to guarantee the protection of that user's privacy. Therefore, the problem with anonymity and Big Data is that it is not so difficult to trace back to an individual's identity when a large amount of data sets are available (Baeza-Yates, 2013). Tene & Polonetsky (2013) discuss this phenomenon in terms of an 'incremental effect': once a bunch of data is linked to an identified individual, it is difficult to disentangle them; any association between this data and a virtual identity breaks the anonymity of the latter.

The right to remain anonymous is strictly connected to the concept of privacy: individuals should have the ability to choose to remain anonymous if they are to preserve the same protection for their privacy online as they currently enjoy offline. Organisations not being able to do so effectively might face the ethical consequences of violating the users' right to privacy.

5i) Individual privacy

While Wright (2011) describes privacy in terms of the four dimensions of *Privacy of personal communications*, *Privacy of the person*, *Privacy of personal behaviour* and *Privacy of personal data*, literature on Big Data do not refer specifically to violations of privacy in only one of such dimensions. Privacy is instead discussed as either *individual privacy* or *group privacy*, where the latter refers to "the protection of information of a group". Big Data analytics uses great amounts of data to infer predictive information about groups of people, and this raises the concern of protecting these groups from potential harm due to invasive and discriminatory data processing (Mantelero, 2017). Potential violations of privacy by Big Data systems concern personal behaviour, sexual and religious preferences, but also the monitoring of communications and the tracking of an individual's movements. In this paragraph, the value of privacy is discussed without giving specific attention to the kind of data is at risk of being identified, coherently with the way that privacy is addressed in existing literature.

Mittelstadt, Allo, Taddeo, Wachter & Floridi (2016) refer to informational privacy as the right of data subjects to shield personal data from third parties. De-individualisation of subjects, discrimination and opaque decision making all constitute threats to the individual's informational privacy. Privacy problems could also result from the activities of abusing the informational reuse and accessing unauthorised data: information reuse involves organisations making new uses of the personal information they have collected, whereas unauthorised access involves employees viewing personal information they are not authorised to view. Both activities can potentially threaten the individual's privacy and result in harms for the individual: for example, they might lead to identity theft or identity fraud (Herschel & Miori, 2017).

The ethical use of data involves knowing how to use data and how to protect privacy and maintain the confidentiality of data. This might translate in removing identifying information from a data record, keeping track of who has access to data, when and how or in knowing the process by which insights are generated.

However, the fact that the processing of Big Data is automated means that the devices that use analysis algorithms are insensitive to privacy issues. And even when humans are involved in the process, the volume of Big Data that they deal with make any effort to protect the individual's informational privacy difficult and impractical (Herschel & Miori, 2017).

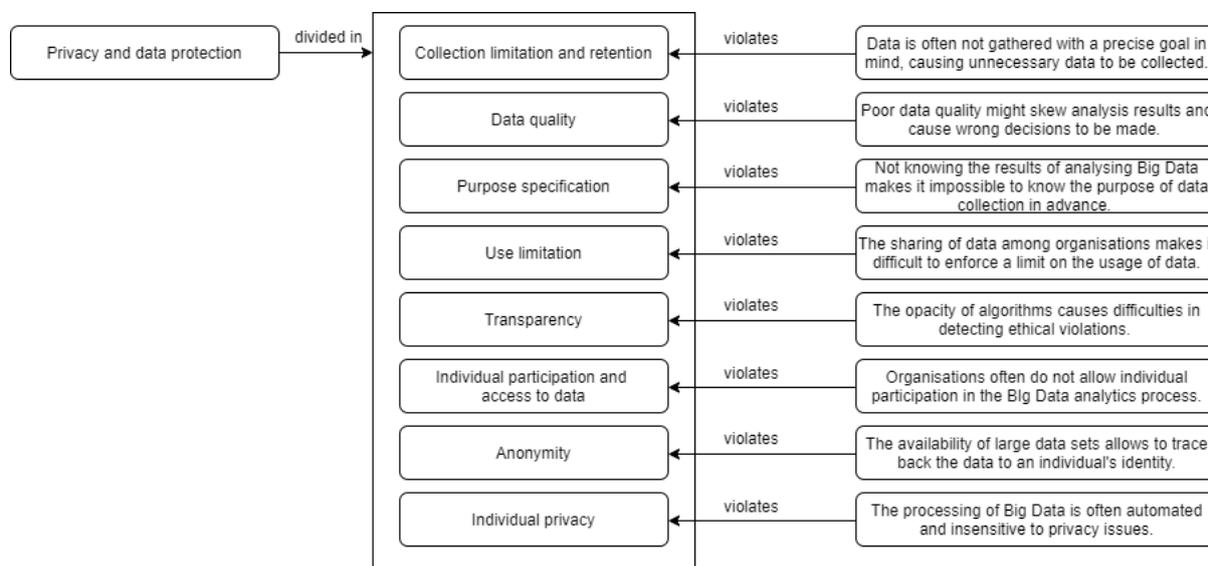


FIGURE 10 – UNETHICAL ACTIVITIES FOR THE PRIVACY AND DATA PROTECTION PRINCIPLE

2.5. Data laws and regulations

Being ethical is not the same as following the law. The law often incorporates ethical standards that most citizens live by. However, laws can deviate from what is ethical. The way laws can deviate from what is ethical is very evident, for example, in the American pre-Civil War slavery laws and in the old apartheid laws of South Africa (Velasquez, Andre, Shanks, J., & Meyer, 2010).

When it comes to information technologies and data, while laws codify some ethical principles, legislations cannot keep up with the risks associated with their evolution. The data environment is evolving rapidly (DAMA International, 2017): we have some privacy rules to govern existing flows of personal information, but we lack rules to govern novel flows, uses and decisions derived from that data (Richards & King, 2014).

Thus, we argue that organisations should work to protect data entrusted to them without waiting for laws to enforce ethical behaviour on them first.

To show how complying to laws does not necessarily mean behaving ethically, we shall consider another example discussed by Cathy O'Neil in her "Weapons of Math Destruction" book. She states that some insurance companies have been using predictive algorithms to determine the risk of any individual to default on a loan are, based on some potentially discriminating parameters. The parameters may include, for example, the zip code of an individual. At first sight this parameter does not seem to be discriminating the individual: chances are that borrowers living in poor areas will default on repaying the loan, thus the algorithm will assign them a low score and target them as a riskier demographic. However, the algorithm is in fact expressing the opinion that the history of human behaviour in that patch of land should determine what kind of loan a person who lives there should get. This phenomenon, besides potentially constituting a violation of that individual's rights, generates a feedback loop that causes those who are already struggling to receive less credit and higher interest rates.

The consequence of companies being able to dive into largely unregulated pools of data to perform predictive analytics is that, by doing so, they can largely avoid government oversight. The danger in using these

algorithms is therefore the fact that they draw from this pool of data and use race and zip code as proxies for financial responsibility: this is unfair and probably illegal and thus raises ethical questions such as “Is it fair to deny a person a loan because they live in a poor neighbourhood?”. While these proxies might work sometimes, what can happen is that a person might be misunderstood and placed in the wrong ‘bucket’. The absence of a feedback loop that will set the system straight makes it impossible to recognise that valuable potential customer has been discarded (O’Niel, 2016).

The following paragraphs will focus on describing the new General Data Protection Regulation of the European Union, including the principles that it is comprised of and the role it occupies in the debate of organisations dealing with data ethics.

GDPR

In the domain of data ethics, the General Data Protection Regulation (GDPR) was issued on May 25 2018 to give EU residents more control over their data. The regulation seeks to reinforce the effective protection of EU citizens in the age of the Internet, as “data travels across the world faster than the time it takes to click” (Granger & Irion, 2018). Besides ensuring that data gets collected legally, the law obliges companies to protect that information and safeguard it from misuse (Matthews, 2019). The regulation serves to replace the fragmented laws and regulations that previously applied on a state level across Europe; thus, it simplifies rules for companies acting within the European market. Furthermore, this legislation “applies to all companies processing the personal data of data subjects residing in the Union, regardless of the company’s location” (Trunomi, n.d.).

The main GDPR statements are illustrated in the table below, as described by Velasquez, Andre, Shanks, J. & Meyer (2010):

GDPR Principle	Description of principle
Breach Notification	Where a data breach is likely to “result in a risk for the rights and freedom of individuals”, a notification of such breach must be done within the first 72 hours of becoming aware of it.
Right to Access	Data subjects have the right to know whether or not personal data concerning them is being processed, where and for what purpose. The data controller shall also provide a copy of the personal data, free of charge, in an electronic format.
Right to be Forgotten	The data subject can demand the erasure of his/her personal data and halt the processing of the data when it is not relevant anymore for the original collection purposes, or when the data withdraws consent.
Data Portability	The data subject has the right to receive the personal data concerning them from a data controller and transfer it to another controller.
Privacy By Design	The requirements of the regulation have to be included and integrated in the design of the data controller’s systems. The data minimisation article also forces the data collector to limit the collection and processing of data to the minimum necessary.
Data Protection Officers (DPO)	The role of a Data Protection Officer (DPO) is a role with expert knowledge on data protection and practices, which is mandatory to have for those controllers whose core activities consist of regularly monitoring data subjects on a large scale.

Granger & Irion (2018) notice that the complexity of the EU data protection law, and practical difficulties in its operationalisation will be major factors undermining its effectiveness. Corporate actors or civil society organisations struggle to be fully compliant with the specification of this law; these difficulties are exacerbated by rapid technological change, which often leave law and legal procedures lagging behind.

Nonetheless, the IBE (2018) states that the GDPR is not just a set of compliance rules. What it does is to highlight the importance of applying ethical values to decision-making, establishing data processes in a transparent way and monitoring employees on the associated behaviours. Non-compliance could mean facing significant fines (up to 4% of the annual global turnover or €20 million), loss of trust, negative publicity and reputational damage. The IBE also argues that compliance should not be however only driven by the fear of consequence; rather, ethics 'start where the law ends'. They suggest that organisations should use the GDPR as a starting point to develop an ethical culture where the importance of the ethical use of personal data is communicated from the top to the employees throughout the entire organisation. Rather than only being prepared for the GDPR to demonstrate compliance with its requirements, the business leaders should develop their strategy for data protection and privacy in an ever-evolving way, in order for it to be able to respond to changes in circumstances (ibe, 2018).

2.6. Data Governance

Organisations have been investing huge amounts of money and efforts on developing their Big Data analytics capabilities, but often do not have a clear understanding of how to use them ethically. These organisations often lack Data Governance practices – that is, defined standards and procedures, for collecting, analysing and using the retrieved insights in an ethical way, as well as training on ethics of employees, ethical leadership and control mechanisms for unethical behaviour (Asadi Someh, Breidbach, & Davern, Ethical Implications of Big Data Analytics, 2016).

In order for organisations to capture value from their IT initiatives such as business intelligence – value that can be in the form of effective decision-making and increased productivity, it is important that the quality of data is being appropriately monitored. Problems of data quality emerge due to the fact that data is spread across disparate systems within an organisation, as well as the fact that data is being collected and used by various levels of an organisation (Cheong & Chang, 2007).

As data keeps growing in quantities, and new technologies allow better, faster, and cheaper storage and processing of such data, organisations face the challenge of developing governance mechanisms that can balance out risks and benefits of Big Data and Big Data analytics. These policies and structures should protect data from the factors that could destroy or limit its value (Tallon, 2013), such as the aforementioned data quality – which does not only affect the ability of the organisation to derive value from data, but also obstructs the compliance to data regulations and ethical norms. A Data Governance program would allow data managers to manage data and its quality as an enterprise asset (Cheong & Chang, 2007).

Defining Data Governance

In order to address data quality issues, organisations should adopt a holistic approach, focusing on “people, processes and technology” and need to constantly quantify and measure their data quality. This implies that in order to address data quality issues, data needs to be governed. Together, people, process and technology allow the creation of a “consistent and proper handling of an organisation’s data across the enterprise” (Wikipedia, n.d.).

A proposed definition of Data Governance that addresses data quality is mentioned in the review of Cheong & Chang (2007) and is the following:

“Data Governance is the process by which a company manages the quantity, consistency, usability, security and availability of data”.

Data Governance defines policies and procedures to ensure proactive and effective data management. The adoption of a Data Governance framework also enables collaboration from various levels of the organisations

to manage enterprise-wide data: this allows the organisation to address issues related to data, such as quality, more easily; furthermore, it provides the ability to align various data related programs with corporate objectives (Cheong & Chang, 2007).

Practitioners such as the Data Management Association (DAMA) define Data Governance as:

“The exercise of authority, control and shared decision making (planning, monitoring and enforcement) over the management of data assets”.

This view goes beyond the idea that Data Governance is only about specifying a framework; rather, Data Governance can also be practiced. From a theoretical standpoint, Data Governance describes processes and defines responsibilities to manage data and information appropriately. In practice, data managers work within this framework by turning their rights (what decisions regarding the handling of data are allowed) and duties (what the related decision-making tasks are) into actions.

Otto (2011) reviews literature on Data Governance, noticing that both the scientific community and practitioners agree on the notion that data is a company asset, the value of which organisations need to maintain and/or increase. Those organisations that establish a formal Data Governance program exercise control over data in an intentional way, and this ultimately allows them to increase the value they get from their data assets (DAMA International, 2017). Besides getting value out of an organisation’s data, the purpose of Data Governance is to ensure that data is managed properly, by establishing how decisions are made about data and how people and processes are expected to behave in relation to data.

(DAMA International, 2017) (Otto, 2011)

An effective governance framework involves four key components:

- 1) *Standards*: data governance establishes standards for data in an enterprise, which can be in the form of data definitions and taxonomies, master data definition, enterprise data models, plus the development and enforcement of technical standards related to data.
- 2) *Policies and processes*: data governance establishes and enforces policies and processes around the creation, development, control, management and audit of data. These can be, for example, in the form of data-related business rules, mechanisms to monitor data and manage changes to it.
- 3) *Organisation*: when launching a data governance initiative, the company needs to address the design of the organisational structure. The initiative will therefore involve the definition of roles and responsibilities within the organisation that are accountable for data, roles that act at different levels and involve both business and IT employees. Examples of these roles are data stewards and data analysts.
- 4) *Technology*: organisations that intend to launch a data governance program should do so only with an underlying technology infrastructure that supports it. Technology can help automate and scale the development and enforcement of data governance standards, policies and processes. Particularly, a data integration technology platform can help automate data-related processes, which involve the access, cleanse, transformation and monitor of data.

(Panijan, 2010)

As a last point, it can be noticed that literature sources often use the terms Data Governance and Data Management in the same context assuming they have the same meaning. While the two terms are used interchangeably, there is a distinction to be made between them. The DAMA makes the difference between Data Governance and Data Management clear by stating that Data Governance oversees Data and Information by ensuring that data is managed properly, while Data Management has execution duties and directly manages data to achieve goals. This distinction is depicted in the figure below:

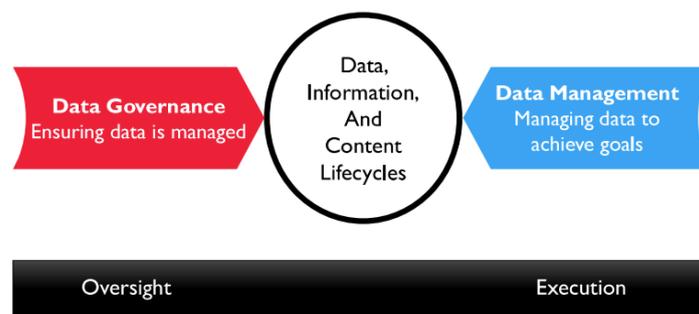


FIGURE 11 – RELATION BETWEEN DATA GOVERNANCE AND DATA MANAGEMENT (DAMA INTERNATIONAL, 2017)

The role of Data Governance

Effective management of Big Data could bring great opportunities for organisations and provide them a way to address the increasing ethical challenges related to Big Data. Big Data ethics raise questions about the ethical nature of business processes; unethical behaviour puts organisational reputation at risk and increases the need of a Data Governance program that could contain the risks emerging from Big Data and its related technologies.

Data Governance practices have a social and legal responsibility to safeguard personal data that, if compromised, would threaten the individual's privacy. Oftentimes organisations use a reactive approach to Data Governance, meaning that only when they find themselves in predicament – facing for example a data loss or privacy breach, they reactively proceed to establish or enhance Data Governance. Such approach is not recommended because data is increasingly strategic and these data-related events could seriously endanger the organisation and put its reputation at risk. For Data Governance to be truly effective, it must be planned with the goal preventing risky data-related events from occurring in the first place (Tallon, 2013).

An ethical approach to data use is increasingly being recognised as a competitive business advantage. Handling data in an ethical way increases the trust of the organisation in its data, process and outcomes of the data lifecycle. Ultimately this can help develop better relationships between the organisation and its stakeholders – who expect ethical behaviour and outcomes from the business and its data processes.

Data Governance intervenes in this context to create an ethical culture, which in practice means introducing controls to ensure that the outcomes of data processing are ethical and do not violate ethical principles, human rights and data regulations. In order for governance to succeed, the organisation as a whole needs to be aware of the risks associated with the misuse of data, as well as be willing to commit to handling data based on principles that safeguard individuals and the company's reputation (DAMA International, 2017).

Organisations that have the ability of unlocking value from their data faster than their competitors will likely be the winners in the race to see who can get the most benefit from Big Data. Data Governance practices will likely reveal themselves to be instrumental in this race (Tallon, 2013).

Organisations protect data based largely on laws and regulatory requirements. However, by holding such big amounts of data, organisations also hold the power of making decisions that affect people's lives. Data managers should therefore recognise that there are ethical, as well as legal, reasons to protect data and ensure that it is not misused. These professional figures have an ethical responsibility to manage data in a way that reduces the risk that it may misrepresent, be misused or be misunderstood. This responsibility extends across the data lifecycle, from the moment data is created to the moment it is destroyed (DAMA International, 2017).

We have defined Data Governance as the process that sets standards and policies that allow organisations to oversee their handling of data, as well as to tackle and potentially improve the quality of their data. From the literature analysis in this research, data quality has been identified as one of the factors that might cause unethical decisions to be made; a Data Governance program should not only be put in place to increase productivity and the effectiveness of decision making, but also to comply to data regulations and ethical principles. Data Governance should be proactively designed within an organisation, making sure that ethics are taken into consideration in such design, with the goal of reducing the risks of unethical data-handling

behaviour and legal consequences due to non-compliance to data regulations.

2.7. Discussion

The presented literature review has attempted to answer the research question: “*What is Big Data ethics?*”. The researcher has led up to a definition of Big Data ethics by first introducing and defining the concepts of ethics and business ethics. The concept of business ethics was deemed relevant due to the scope of the research to commercial organisations and to the influence that an ethical organisational culture has on the way that data ethics are dealt with. Following, the concept of Big Data was introduced and information regarding the supporting technologies of Artificial Intelligence/Machine Learning, as well as Data Mining was included in the literature to provide more context to the concept of Big Data itself. Furthermore, literature on data value chain and lifecycle models was reviewed to provide an overview and explain the main activities involved in the process of turning raw data into insights for decision-makers. Existing literature showed that the concept of data ethics hadn’t been discussed before in light of the phenomenon of Big Data. Thus, the researcher combined existing definitions of data ethics with previously determined knowledge on ethics and business ethics to derive the following definition of Big Data ethics, which constitutes the answer to the abovementioned question:

“Big Data ethics is a branch of business ethics that studies ethical problems that arise in the business environment when using Big Data and algorithms for data analysis. Its goal is to develop moral rules, standards, or practices that support moral decision-making based on Big Data analytics.”

This review also tried to answer the question: “*What is the role of existing data laws and regulations in addressing Big Data ethics?*”. The relation between laws and ethics has been shown for both business ethics and data ethics. Starting from business ethics, research has shown that oftentimes businesses try to resist the imposition of regulations on their activities and that, even if organisations didn’t do so, some unethical business practices simply cannot be eliminated by means of regulations e.g. because they would bring more costs than benefits or because they would be too difficult to monitor. Furthermore, when it comes to information technologies, the technological advancements happen so quickly that regulations struggle to keep up with them.

For what concerns data ethics, the new GDPR regulation was discussed: this has been introduced to give users more control over their data. It also has brought within the business environment the assumption that, in order for the business to use data ethically, it is enough to be compliant to governmental data regulations in act such as the GDPR. This review has served to determine that, also in the field of Big Data, being compliant does not necessarily mean being ethical. Thus, being compliant to regulations such as the GDPR does not mean living by foundational ethical values – assuming that living by such ethical values could give an organisation a solid assurance that they are behaving ethically. In favour of this argument, examples have been shown of situations where unethical behaviours occurred that overcame laws. Instead of focusing solely on demonstrating compliance with its requirements, organisations should instead use GDPR as a starting point to develop an ethical culture that would allow them to be able to respond to changes in circumstances (e.g. new data regulations and new technologies).

Further, the literature review served to answer the research question: “*How do fundamental ethical principles relate to the Big Data context?*”. In order to do so, first the ethical principles of Autonomy, Beneficence, Maleficence, Justice and Privacy and Data protection were described, drawing from literature on the ethics of biomedical science, as well as Information Technology ethics: these are the generic principles that should drive a company’s ethical behaviour. For each principle, a list of values/issues strictly related to it were defined; these served to simplify the activity of applying the fundamental principles to the context of Big Data. Establishing a clear relation between the principles and Big Data allows the researcher to determine, in the specific context of Big Data, what can be considered ethical behaviour and what not. Thus, for each sub-value, a concrete example of how it might be violated by an organisation that makes use of Big Data to retrieve insights was shown. These examples were derived from existing papers, which discussed ethical concerns

related to the Big Data phenomenon; previous research was however lacking an explicit mention of the ethical principles on the basis of which they had determined where ethical violations were occurring in the Big Data lifecycle. The researcher thus enriched the existing body of literature by making the relationship between ethical principles and the big Data context explicit, resulting in a list of unethical activities related to the use of Big Data within an organisational environment. A graphical overview of this analysis is depicted in the figure below:

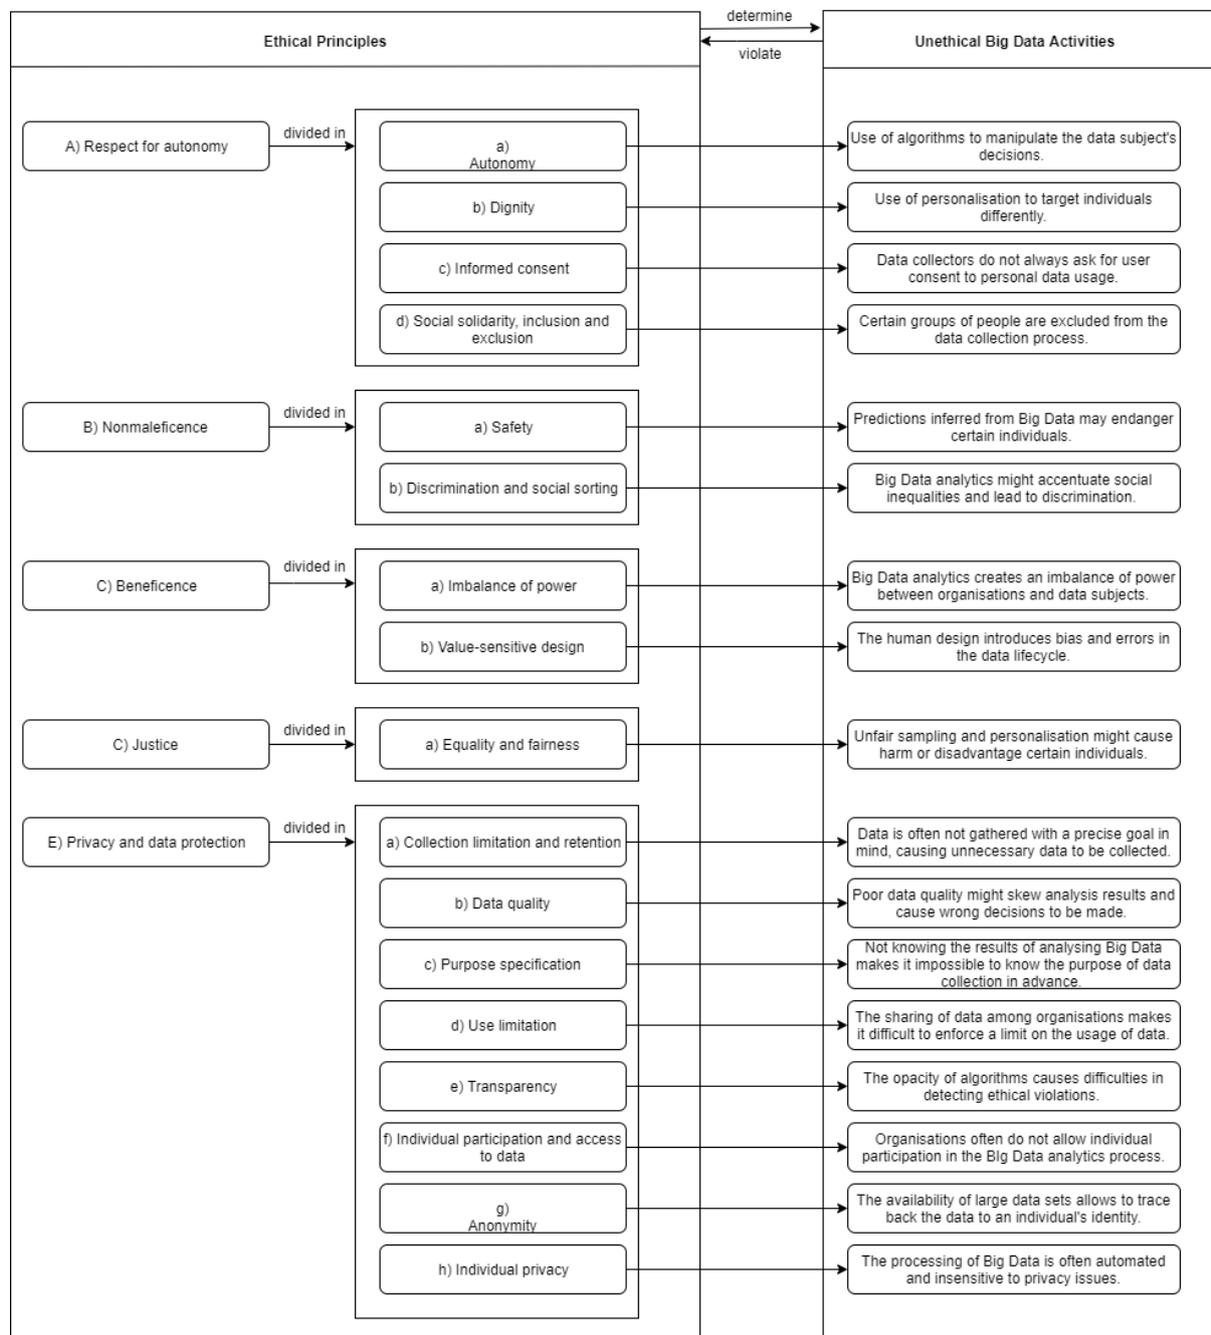


FIGURE 12 – LITERATURE REVIEW FRAMEWORK

Lastly, this review served to describe what Data Governance is and what its role is in addressing the ethics of Big Data. Data Governance was introduced as the process that defines policies and procedures which allow a company to manage data in all its defining characteristics (quality, availability, consistency etc.). It was found that by means of effective Data Governance it should be possible to address the questions raised by Big Data ethics, and contain the risks emerging from an unethical handling of data. Designing a Data Governance program in a proactive way, and taking ethics into consideration while doing so, will enable a firm to prevent

risky data-related events from occurring in the first place, as well as legal consequences due to non-compliance to data regulations.

3. Research Methodology

3.1. Conceptual Framework

The literature review in the previous chapter helped define the major concepts in play in the research. Figure 13 illustrates an extension to the conceptual framework of Figure 2, which summarizes the findings of the literature review by making the relationships between the research concepts more specific; furthermore, the framework makes explicit which research question address which relationship.

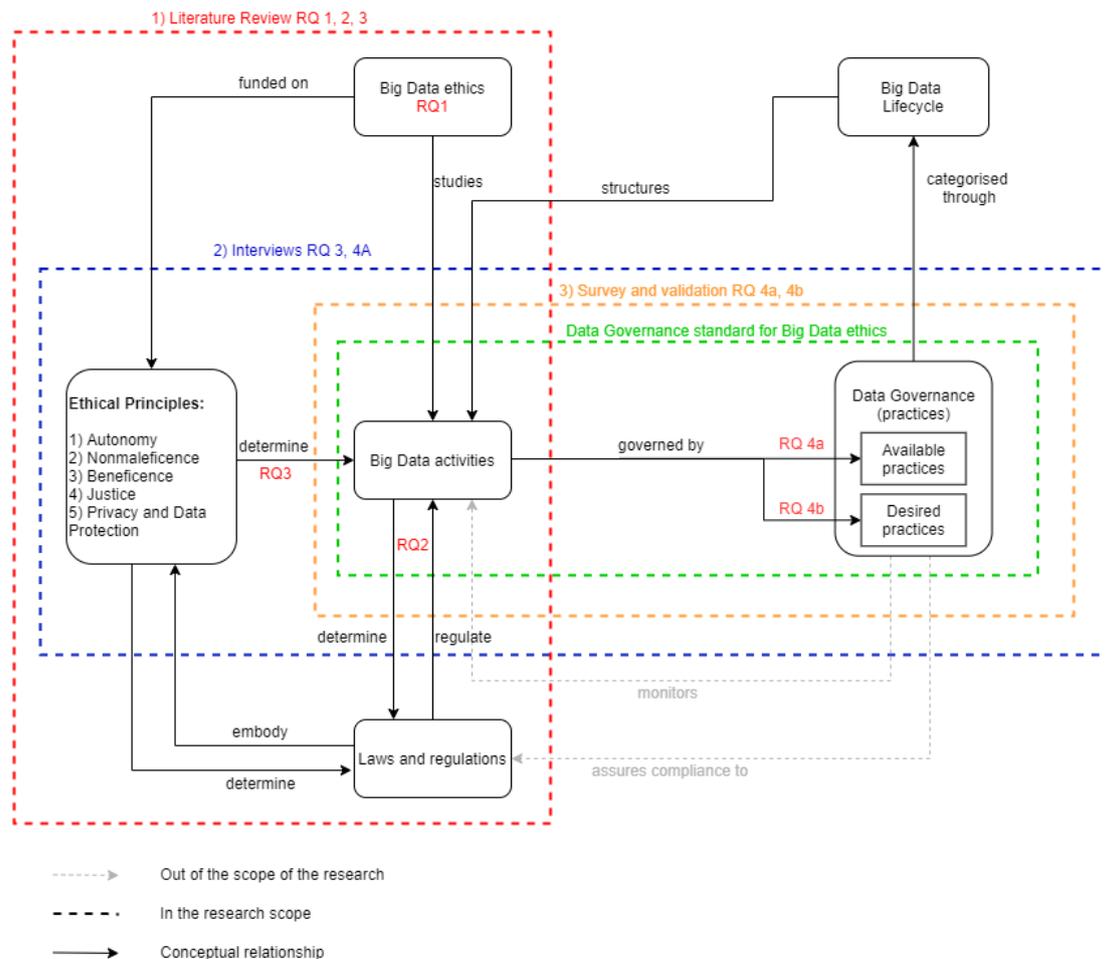


FIGURE 13 – CONCEPTUAL FRAMEWORK

The main subject of the research is Big Data ethics. This has been identified as a field that studies ethical problems arising from the execution of certain activities involving the use of Big Data within organisations, by means of high ethical principles which judge these activities to determine where unethical behaviour occurs. The definition of Big Data ethics was the answer to the research question 1. The application of ethical principles to the context of Big Data contributed to answering the research question 3, and the consequent identification of unethical Big Data activities currently occurring within commercial organisations. The role of laws and regulations in the data ethics debate was analysed through the Research Question 2: laws and regulations have the objective to regulate the use of Big Data within organisations; however, due to the fact that technologies are constantly evolving, it is possible that emerging Big Data activities may not be yet regulated by laws and regulations: the feedback loop in the conceptual framework shows that the emergence of new activities create the need for new laws that regulate them. Research questions 1, 2 and 3 have been

answered in the previous chapter through a review of the existing literature. The relationship between the Big Data activities and the ethical principles that helped identifying them will be validated in a round of interviews with experts.

Data Governance practices were identified as a control factor for the Big Data activities: the activities can be governed by means of governance practices, which can in turn support organisations into becoming more ethical bodies. Also, Data Governance through the definition of policies and procedures monitors the Big Data activities executed within an organisation. Data Governance, in the process of monitoring data activities, ensures that these are compliant with existing regulations. Governance practices for handling data ethically are investigated in the research questions 4a and 4b. Such research questions distinguish between the practices that are desirable to implement within commercial organizations – which will be investigated through the expert interviews, and the practices that are currently in place within such organizations – which are addressed by the survey executed as part of this research.

Once exact relationships are identified between unethical Big Data activities and both desirable and used Data Governance practices, these will be illustrated in the deliverable of the research, namely the Data Governance standard for Big Data ethics.

The Big Data lifecycle was introduced as an instrument that explains the placement of specific Big Data activities within the data processes of organisations and helps categorise the identified Data Governance practices in a structured way. It will also be used to give structure to the survey and the way questions are asked to respondents, as well as to the Data Governance standard.

3.2. Methodology approach

For this study, an inductive research approach is used: such method of reasoning consists in working ‘bottom-up’, “using the participants’ views to build broader themes and generate a theory interconnecting the themes” (Soiferman, 2010). This approach allows the researcher to build upon the broad themes of ethics, Big Data and Data Governance – which have already been subject, individually, of research, and find the relationships that interconnect them to generate a novel theory.

The inductive approach is associated with a qualitative type of analysis: in fact, qualitative research employs induction reasoning in that “it moves from specific observations about individual occurrences to broader generalisations and theories” (Soiferman, 2010). The researcher adopting this approach qualitatively collects data by gathering the words of the participants, and then moves onto analysing them by detecting common themes and patterns in the data.



FIGURE 14 – INDUCTIVE METHODOLOGY

A qualitative type of analysis requires the researcher to subjectively interpret the words of participants, thus potentially introducing bias while doing so. However, one of the advantages of qualitative research is being able to form a tentative, early hypothesis from the identified themes, potentially leading to inductively developed theories or general conclusions. Furthermore, qualitative research methods are recommended as a method of collecting data about people’s subjective experience, their views and perceptions (Burnard, 1999): given the controversies around the topic of data ethics, and the multitude of perspectives surrounding it, qualitative methods are used in this study to allow its complexity to emerge from the participants’ perspectives (Soiferman, 2010).

The literature review is used to provide evidence for the purpose of the study and to identify the underlying problem that will be addressed by the researcher (Soiferman, 2010), namely the problem of commercial companies using Big Data unethically to make decisions that might negatively impact the end users. Once the hypotheses and research questions are narrowed, data is gathered through interviews with experts: open-ended questions are asked to learn from the participants' experiences, allowing the exploration of a variety of points of view (Burnard, 1999). The collected data is then analysed qualitatively: due to the lack of structure of the textual data retrieved from the interviews, the researcher has to provide his own interpretation to it, and organise it into themes and categories to generate a consolidated picture from it. Such interpretations are shaped by the personal stance of the researcher, which depends on his experiences and backgrounds (Soiferman, 2010). Due to the possibility of introducing bias in the qualitative analysis of data, it is important that the interpretations made by the researcher are supported by a stage of validation. Various sources are used to verify the theme of the research (Soiferman, 2010): validation is executed in several stages of this study, sometimes in parallel with data collection, and utilising different methods, namely a survey and validation sessions.

3.3. Methodology overview

The research can be seen as a 4-step process, which is illustrated in Figure 3. *Phase 1* consists in a systematic literature review and aims to answer the research questions 1, 2 and 3. During this phase, the concept of Big Data ethics is defined, building up on existing definitions of Ethics and Business Ethics found in literature. Furthermore, the role of laws and regulations in addressing Big Data ethics is addressed in this phase. Lastly, fundamental ethical principles are defined by means of literature: these have been historically used to define ethics in the biomedical field, but in more recent years they have been also applied to information technology and AI. In this study, these principles have been applied to the context of Big Data: establishing this connection is important because the ethical principles help conduct an ethical assessment of Big Data activities.

Phase 2 consists in a round of expert interviews which serves to validate the findings of the literature review, and specifically the relationships identified between fundamental ethical principles and the context of Big Data. Furthermore, this phase serves to collect information regarding the Data Governance practices that can be used to address ethical problems concerning the use of Big Data within commercial organisations. By relating the Big Data activities to the governance practices, it is possible to construct a standard of Data Governance practices to address the problem of data ethics within an enterprise.

Phase 3 consists in building and conducting of a survey for collecting further information regarding the Data Governance practices that would be desirable to use to tackle unethical Big Data activities within commercial companies. Within these practices, the survey aims to collect information regarding which are currently in place in the landscape of commercial organisations, which together creates an overview of the status of ethical treatment of data within the surveyed companies.

Lastly, *Phase 4* is a validation round executed by means of expert interviews, in which the information collected in the survey is reviewed and validated, in order to generate a validated version of the governance standard. Together with Phase 3, Phase 4 contributes to answering the research question 4, as well as the main research question.

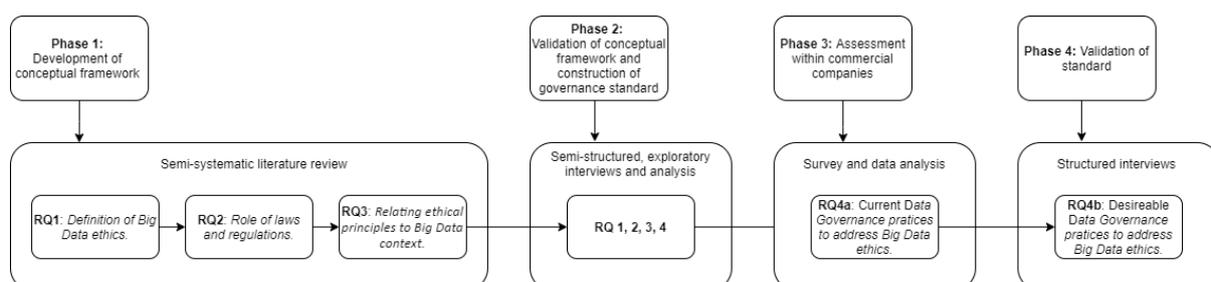


FIGURE 15 - METHODOLOGY PROCESS

The following sub-chapters describe each phase more in depth.

3.4. Literature review methodology

A semi-systematic methodology was chosen to conduct the literature review: a systematic literature review provides a structure to the way papers are selected, processed and turned into outcomes. A good systematic literature review might achieve the following:

- Establish to what extent existing research has progressed towards clarifying a particular problem (Siddaway, 2014): for this research, a systematic kind of literature review proves to be a useful tool, due to the lack of a standardised solution for addressing Big Data ethics within organisations.
- Identify relations and gaps in the literature (Siddaway, 2014): the lack of a comprehensive definition of the concept of Big Data ethics was identified in existing literature; furthermore, even though existing literature has identified Data Governance as a method to address ethical problems concerning the use of Big Data, a concrete description of practices to do so is lacking. A systematic literature review methodology is also useful to define the concept of Big Data ethics, which currently lacks an agreed-upon definition in literature: this is done by identifying relationships between the topics of ethics, business ethics and Big Data.
- Comment on, evaluate, extend or develop theory (Siddaway, 2014): the systematic literature review can be used as a starting point to develop a data governance 'theory' of practices, to be used to address problems concerning the use of Big Data for decision making in commercial organisations – problems which are already known in existing literature.

The list above provides motivations to pursue a systematic kind of literature review. However, the researcher has chosen not to adhere too strictly to the rules of a systematic literature review. A traditional systematic review can be considered a research in its own right (Siddaway, 2014), however in this study the literature review was only to address particular research questions.

Traditional literature reviews and systematic literature reviews share the goal of identifying published literature on a topic, critically appraising, and summarising the critical points of the current knowledge about a problem (JEPS Bulletin, 2018).

The literature review in this research contributes to answering the following literature review questions:

[1] *“What are Big Data ethics?”*

[2] *“What is the role of existing data laws and regulations in addressing Big Data ethics?”*

[3] *“How do fundamental ethical principles relate to the Big Data context?”*

The starting point for the execution of the literature review is the clarification a priori of the objectives of the review, which are based on the previously defined research questions (JEPS Bulletin, 2018). The following goals for the literature review were identified:

- RQ [1]: The definition of the concept of Big Data ethics, in such a way that it is rooted in literature.
- RQ [2]: The clarification of the role that laws and regulations plays in enforcing ethical behaviour from a data perspective within organisations.
- RQ [3]: The application of existing theory of ethics and ethical principles to the specific context of the research, in order to determine how the concepts of Big Data and ethics relate to each other in practice.

The second step involves the breaking of the research questions down to individual concepts to create search terms: these terms effectively operationalise the research questions to find as many relevant articles as possible to include in the review (Siddaway, 2014). The keywords chosen for each research question are

summarised in Annex A. The search of papers is conducted on electronic databases, using first both Google and Google scholar; further papers are found on Science Direct, Research Gate, IEEE Xplore, Semantic Scholar and Springer. The snowballing technique is partially executed to find relevant material within the reference lists of the so-found papers. The results are sorted by relevance; no filtering of publication year is executed: however, the majority of the papers found have been published after the year 2011, with some exceptions of papers which were written between the years 1985 and 2000, showing that the grand majority of the body of literature on the selected subjects has only been researched in the last two decades.

The third step consists of setting the exclusion criteria which apply for the selection of the most relevant papers: these are the criteria that the literature material must meet in order to be excluded from the study (JEPS Bulletin, 2018). The exclusion criteria used are different based on the research question and are the following:

- RQ [1]: Papers approaching the topic of Big Data ethics from a too narrow scope (ethics of Big Data that affect a specific area/sector) were excluded. On the same line of reasoning, papers considering the ethics of data-driven research were excluded as well due to their focus on research data. Paper discussing the ethics of Big Data from a security perspective were also excluded due to the fact that security is out of the scope of the research.
- RQ [2]: Sources discussing data legislation in act outside the European Union were excluded due to the European focus of the research.
- RQ [3]: Papers discussing the ethical concerns around AI were excluded due to their strict technological focus. Papers discussing the possibility of turning machines into ethical agents (also referred to as Machine Ethics) were excluded because not relevant for this study.

Based on such criteria, the abstracts and conclusions of each paper are analysed to make the exclusion and to select the most relevant papers.

The fourth step consists of the creation of a clear record keeping system that allows a systematic organisation of the found papers (Siddaway, 2014). The literature sources are inserted and catalogued in an Excel sheet to keep track of which papers need to be fully reviewed, and of which ones have been excluded and why. This system ultimately allows to record what the researcher has done and the decision-making involved in the process. The catalogue of papers follows the structure illustrated in the figure below:

Authors	Year	Name	Link	Selection	Status	Keywords	Notes	RQ
White & Ariyachandra	2016	Big Data and Ethics: Examining the Grey Areas of Big Data Analytics	https://www.academia.edu/29739493/BIG_DATA_AND_ETHICS_EX	Good	Reviewed	Information Technology (IT), Ethics, IT and Ethics, Analytics and Privacy, Security and Big Data Analytics	Used for identifying the ethical challenges of privacy, security, ownership and evidence-based decision making.	Introduction + 3 (Ethical Concerns)
Leonelli	2016	Locating ethics in data science: responsibility and accountability in global and distributed knowledge production systems	https://www.researchgate.net/publication/310822987_Locating_ethics_in_data_science	Excluded	Reviewed	Big data, epistemology, knowledge production, ethics, science policy, research governance	Excluded because it discusses the ethics of data management in the research environment.	None

FIGURE 16 - RECORD KEEPING SYSTEM OF LITERATURE SOURCES

The fifth and last step consists in deciding whether a qualitative or quantitative research synthesis is most appropriate. A qualitative type of synthesis is chosen due to the fact that the studies involved in the review are methodologically diverse, which makes a quantitative analysis impractical (Siddaway, 2014). Furthermore, a qualitative research synthesis is deemed appropriate when developing a new theory (Siddaway, 2014), which in this case concerns the development of a novel Data Governance standard for Big Data ethics. At this point a synthesis of the results can be written, by working on the selected literature sources to integrate their findings and interpret them in a narrative form (JEPS Bulletin, 2018).

In order to achieve the objective of the research question 1, the underlying topics behind Big Data ethics were analysed in isolation, and then gathered: definitions of ethics and business ethics were provided to build up to a comprehensive definition of Big Data ethics, aided by a few sources that previously defined the concept of data ethics.

In order to achieve the objective of the research question 2, the concept of the law was contrasted to the concept of ethics to determine which place regulations occupy in the Big Data ethics debate. Furthermore, existing regulations active within the European Union were investigated to create a picture of what is currently in place to guide ethical data behaviour within commercial companies.

In order to achieve the purpose of the research question 3, ethical values are introduced based on existing sources in literature. These fundamental values are then used as a guideline and applied to the context of Big Data: unethical activities related to the use of Big Data for decision making purposes are searched for in existing literature and categorised – so that they can be each associated with an ethical value. The Big Data lifecycle is also defined to provide structure to the Big Data ethics debate and to be able to understand where each Big Data activity is positioned within the data process of a company.

The results of the literature review resulted in a framework of ethical areas of interest for the ethics of Big Data, as well as sub-aspects that refer to each of these areas.

An overview of the literature review method above described is illustrated in the figure below.

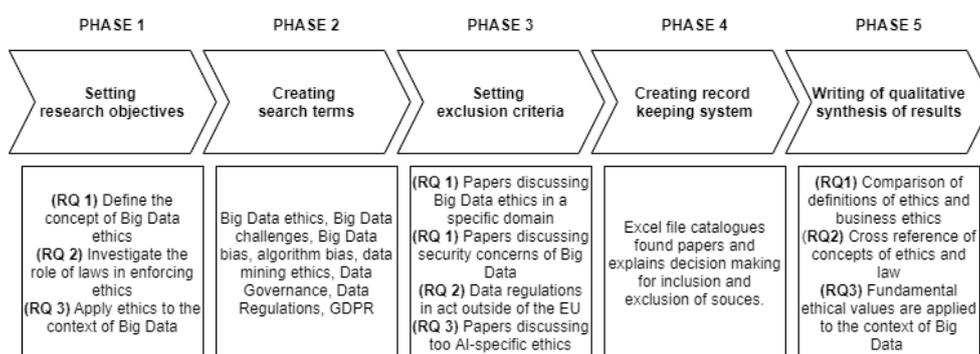


FIGURE 17 - METHODOLOGY OF LITERATURE REVIEW

3.5. Interview methodology

By means of the literature review, the researcher was able to construct a framework which depicts fundamental ethical principles and how these relate to the context of Big Data. In doing so, the researcher implicitly conducted an ethical assessment of the Big Data activities mentioned in existing papers, which then led to a list of relationships identified between the ethical principles and the context of Big Data; due to the way the framework was built, it requires validation to determine whether the framework is complete and whether there are any activities missing. Furthermore, very little material on the Data Governance practices needed to address unethical Big Data activities was identified in literature, thus the need to investigate this topic further. An overview of these objectives is shown in the table below:

Objectives	Research Question
Validate unethical Big Data activities discovered in literature	3
Investigate previously unidentified unethical Big Data activities to integrate in framework	3
Appoint the critical ethical issues – associated with a specific phase of the Big Data lifecycle, and investigate potential solutions to them in the form of Data Governance practices	4a

The chosen method of data collection to achieve the described objectives is semi-structured interviews. An interview is a conversation whose purpose is to gather descriptions of the world of the interviewee and how they interpret the meanings of the phenomena considered in the research (Alshenqeti, 2014). Interviews are generally used in conducting qualitative research and have the potential to generate useful insights into the opinions and experiences of the interviewees, as well as understandings of processes and behaviours. Interviews are especially a useful instrument when it is possible to identify some people who are in key positions to understand a situation: more details and insights can be collected by conducting an interview with

them, rather than through a questionnaire (Rowley, 2012). For this research interviews were considered an appropriate method of data collection due to the possibility to identify specific expert figures, able to give insights on the subject of Big Data ethics. Also, as interviews are interactive, the interviewee is able to probe into emerging topics (Alshenqeeti, 2014), as the one of Big Data ethics. Furthermore, ethics is a subjective topic with no absolute truths and the setting of an open discussion may be more appropriate to ensure mutual understanding between the researcher and the data collection source: during an interview, the interviewer may rephrase or simplify questions that were not understood by the interviewees, to guarantee a more accurate collection of data. Also, since interviews can be recorded, an accurate report of the insights provided by the interviewees can be produced (Alshenqeeti, 2014).

Interviewees can be executed in different ways, but for this research the choice fell on semi-structured interviews due to the flexibility that it gives the interviewer to probe and expand the interviewee's responses (Alshenqeeti, 2014).

When executing semi-structured interviews, it is important that some structure is given to the interviews, which guarantees that all relevant areas are discussed and thus that the answers given by the interviewees contribute to answering the research questions – while giving the interviewer the freedom to ask follow up questions and going more in depth on certain answers during the interview itself. This structure comes in the form of an Interview Protocol, the design of which is executed based on the guidelines laid down by Castillo-Montoya (2016) in his Interview Protocol Refinement Framework (IPR framework) depicted in the figure below.

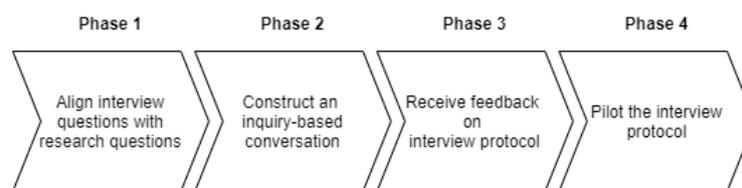


FIGURE 18 – CONSTRUCTION OF THE INTERVIEW PROTOCOL ADAPTED FROM (CASTILLO-MONTOYA, 2016)

Following the first phase of the framework, an initial Interview Protocol is drafted, which includes questions divided in three areas:

- 1) Definition of Big Data ethics, including questions that set the scene of the interview
- 2) Ethical principles applied to Big Data, including questions that validate the literature review framework
- 3) Big Data ethics and laws/regulations, including questions that provide more context on the Big Data ethics debate, aimed at discussing the role that laws and regulations play in enforcing ethical usage of data within organisations

During the second phase the Interview Protocol is transformed into an inquiry-based conversation with the addition of follow-up questions. Furthermore, the formulation of the questions is reviewed to make sure they are written differently than the research questions. Transition questions are added to ensure that the key questions are not asked directly, but that the conversation naturally leads up to them. Lastly, the opportunity to provide feedback at the end of the interview is ensured through a dedicated section of the Interview Protocol.

The third phase is dedicated to receiving feedback about the developed protocol. The protocol was evaluated by two Data Management experts, who checked whether the questions were formulated in a correct and clear language; the review also led to certain questions being removed from the protocol. Lastly, the fourth phase consists of piloting the Interview Protocol: the first interview served as pilot, but due to resource constraints the results of the interview is also used in the study analysed in a successive stage of the research. The pilot of the Interview Protocol served to show that the interview duration was too long: as a consequence, the contextual questions of section 3 were removed and the questions of section 2 were shortened. The pilot also served to understand the need to investigate Data Governance practices more in depth. The changes made in the third and fourth phases of the IP Refinement framework lead to the following Interview Protocol sections:

- 1) Introduction
- 2) Definition of Big Data ethics
- 3) Ethical Principles applied to Big Data
- 4) Critical Big Data activities
- 5) Feedback

The final version of the semi-structured Interview Protocol is visible in Annex B. The first section of the interview allows the interviewee to introduce himself and his position within the company he works for. The second section serves to align the interviewer and interviewee knowledge on the topic of Big Data ethics by introducing the definition of this concept used in the study. Furthermore, some contextual questions are asked on the importance of addressing ethical concerns related to the use of Big Data within organisations: these questions serve as a lead up for the third section of the interview, in which the interviewer goes through the literature review framework and covers all ethical principles and sub categories of those principles, as well as the related Big Data activities; the interviewee is asked to review the described relationships and activities and provide their input based on their personal experience. In the newly added section 4 the interviewee is given the chance to point out the most high risk phases of the Big Data lifecycle in terms of ethics, as well as give some insights on the Data Governance practices that can possibly be used to address the related ethical problems. Lastly, the interviewee is given the chance to leave feedback regarding how the interview was conducted.

A group of 8 experts was identified to take part in the round of interviews: these are people who have knowledge on the Big Data process and are able to perform an ethical evaluation of Big Data activities traditionally executed within commercial companies. Due to the novelty of the field of Big Data ethics, experts from the mixed fields of Big Data analytics and information management, digital ethics, digital law and data privacy were selected to participate in the interview phase of the research. The variety of expertise among interviewees allowed data to be collected from multiple perspectives, which is fundamental when researching the topic of ethics, since those who work directly with data may have a different opinion than policy makers and those who deal with data privacy issues at a managerial level within a company.

The interviewees were asked prior to the interview the permission to be recorded: all participants accepted, thus the eight interviews were recorded and in a second moment transcribed. When sensitive information was disclosed, it was inserted in the transcript but, as requested, not included in the interview analysis.

The interviews were transcribed with the aid of the software f4transkript, and then transported onto the complementary software f4analyse. A systematic method was used to analyse the textual data of the transcripts, following the guidelines of Burnard (1999), who described the process of organising unstructured text by breaking it down into meaning units, developing a category system, and grouping together ideas of a similar sort (Burnard, 1999). The first step consists in cleaning the text, by removing the material that does not relate directly to the topic, or that is peripheral: as a consequence, the introductory parts of the interviews, as well as sensitive information that could not be used in the analysis were stripped out of the transcripts. The second step consists in dividing the text into meaning units: a discrete phrase, sentence or series of sentences which convey one idea (Burnard, 1999). This step was quite intuitive to execute due to the semi-structured nature of the interviews, and it was facilitated by the software feature of notes, which allowed to easily gather related statements together. The ultimate goal of the interview analysis is to identify patterns within the textual data, similarities and differences in the responses of the interviewed experts (Burnard, 1999). In order to do so, the meaning units can be grouped together under common themes, or labels. In this analysis literal categories labels are used to organise the data: these identify in a very literal sense the contents of the interviews (Burnard, 1999). The labels chosen by the researcher correspond to the categories of the literature framework which was discussed during the interviews. The meaning units were gathered under five main labels which correspond to the major ethical principles in the framework. Then, under each of these labels, the meaning units were divided in one *agree* category (including all those interviewee statements that agree with the examples of the framework provided by the researcher) and one *disagree* category (including the statements that disagree with the examples of the framework). The complete list of labels and 'sub-labels' is displayed in Annex C. Lastly, as an additional fourth step, the researcher decided to use the notes feature of

the software to insert comments on the meaning units, such as interpretations of what the interviewee said. An example is shown in Annex D.

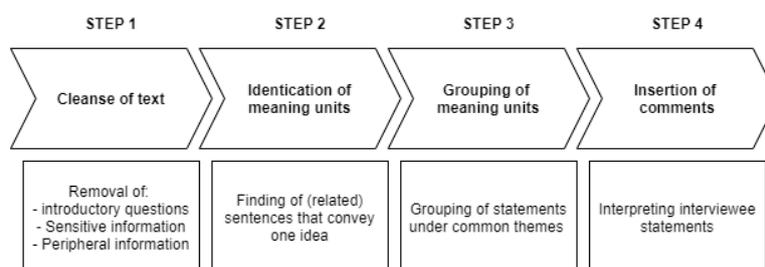


FIGURE 19 – METHODOLOGY OF INTERVIEW ANALYSIS ADAPTED FROM (BURNARD, 1999)

The process of structuring the textual data of the transcripts (Figure 19) facilitated the analysis of the results that followed, in that the categorised text could be transferred onto Excel sheets and be processed more easily on that environment. A framework of critical Big Data activities and corresponding governance practices to address them was constructed as the result of the interviews. More in detail information about the results of the interviews follows in the next chapter of the Interview Analysis.

3.6. Survey construction

At the point of collection and analysis of the interview data, the framework derived from the literature review is validated and research questions 1, 2 and 3 are answered. The interviews also served to determine which Big Data activities are deemed to be critical to address, in order to guarantee ethical behaviour – and the corresponding lifecycle phases they are associated to. Lastly, during the interviews potential solutions for such critical activities were discussed. At this point in the research process, the researcher asked himself whether the governance framework for ethical data usage – constructed by means of literature and the interviews – is complete: that is, are there any statements missing from the framework? Furthermore, research question 4b, which asks whether any of the practices included in the framework are in actual use within commercial companies, still needs to be answered. These goals are depicted in the table below:

Objectives	Research Question
Test governance framework for completeness property	4a
Investigate the Data Governance practices in use within commercial companies to address the ethical concerns raised by the use of Big Data	4b

In order to achieve these objectives, a survey methodology was chosen, specifically a questionnaire. Questionnaires are conducted to gather large size of information in a short period of time (Denscombe, 2010): for this reason, it is deemed to be a good method choice to both test the framework – by collecting a multitude of expert opinions who can fill in the ‘missing pieces of the puzzle’ – and create a comprehensive view of the ethical integrity of commercial companies’ Big Data usage. Furthermore, one of the advantages of the questionnaire methodology is that the members of the sample group can remain anonymous (Denscombe, 2010): given the sensitivity of the ethics topic, presumably companies would not want to disclose private information about the way they handle data, if this is not exactly ethical; the questionnaire allows them to execute an ethical assessment of their data usage that the researcher can also benefit from, without having to publicly unveil such information. Granting respondents the anonymity of results can therefore help make up for one of the disadvantages of the survey methodology: the potential inability of respondents to provide sensitive information to the researcher. Also, given the diversity of roles within an organisation that can deal more or less directly with data, the survey method is able to pull feedback from such a diverse pool of respondents (Foley, 2018), and thus provide an holistic overview of ethical data usage within commercial companies. Given the descriptive nature of the Big Data ethics topic, the survey is used in this research as a qualitative method of data collection, meaning that the questionnaire will include open ended questions that allow more conversational answers (Foley, 2018).

The process of constructing the survey followed the general guidelines of Hensley (1999) with some adaptations, given the qualitative nature of the survey. The procedure is illustrated in the figure below:

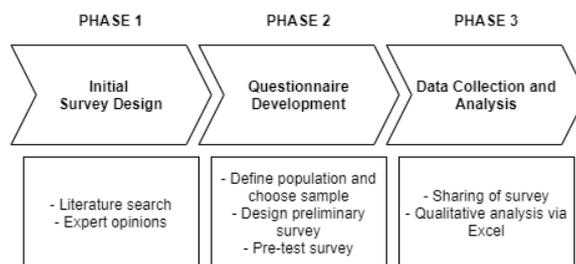


FIGURE 20 – METHODOLOGY OF SURVEY CONSTRUCTION ADAPTED FROM (HENSLEY, 1999)

Hensley describes the first step of the process as a combination of the use of literature sources and expert opinions for the initial survey design. In this study, the initial survey design first consists in the formalisation of the expert opinions collected in the previous interview phase: these statements, together with literature sources, formed the informational basis needed to design the survey.

In the second phase, a preliminary survey is constructed: the tool used to set up the questionnaire was Qualtrics, an online tool that allows the easy distribution of the questionnaire to respondents, as well as the monitoring of responses and the conversion of the ladder in Excel files. Furthermore, in this phase the survey is pre-tested on two Data Management experts and two students that checked the questionnaire for the clarity of questions and answers, as well as its intuitiveness of use. Based on the received feedback, the questionnaire is shortened significantly so that it wouldn't cross the 20 minutes mark, some questions are reformulated so that they cannot be misunderstood, and order bias is eliminated where possible. The final survey consists of five blocks:

- 1) *Introduction to survey*: the research topic is explained, as well as the purpose of the questionnaire. The respondent is given information regarding the anonymity clause of the questionnaire and the expected duration to complete.
- 2) *Background questions*: this block contains personal questions regarding the background of the respondent and the industry they work in.
- 3) *Ethical evaluation*: the concept of Big Data ethics is introduced, and the Big Data lifecycle description used in the study is shown to the respondent. Following, 6 sub-blocks go each in depth about the highest priority lifecycle phases (as concluded from the interviews).
- 4) *Contextual questions*: this block contains questions regarding the context of the research, and more specifically the perceived riskiness of using Big Data for decision making purposes.
- 5) *Feedback*: the respondent is given the chance to leave feedback on the survey.

In each block, a combination of open and multiple choice questions are asked. Please refer to Annex E for the complete survey.

The second phase also consists in the definition of the population and choice of the sample size: for this study, the population is defined on the basis of the three characteristics of Job Title of the respondent, Industry that the respondent works in and Field of Expertise of the respondent. The complete list of characteristics that the ideal respondent should possess is shown in Annex F: the compliance to any one characteristic mentioned in the requirements list is sufficient to consider the respondent capable of filling out the survey. The chosen sample size is of 30 respondents: this number takes into consideration the fact that the pool of potential respondents who are willing to take the time to fill out the survey is quite limited, as well as the fact that diversity of the target population is a more ideal target to reach than the size of the population: it is in fact more efficient to purposely select a diversity sample in a qualitative study (Harrie, 2010) rather than a large, non-diverse sample.

Lastly, the third phase consists in collecting the data and analysing it. In order for the questionnaire to reach enough respondents, the survey was shared with several respondents who matched the requirements via LinkedIn, the LinkedIn group of Leiden University, as well as the personal networks of thesis supervisors and Deloitte consultants (in the areas of Enterprise Data Management, Analytics and Information Management

and Risk Advisory). The message that was sent to the potential respondents is visible in Annex G. Once the target number of responses was reached, the results were converted to an Excel file, where the quantitative analysis of the results was later executed. The exact process for the analysis of the survey results follows in the Survey Analysis chapter.

3.7. Results validation

A governance framework is constructed using data collected through a round of expert interviews and a survey addressing data practitioners. Due to the qualitative nature of such data, the framework may be affected by biases such as misinterpretation of the study participants' opinions from the researcher. A round of validation is therefore required to make up for potential biases introduced in the framework, as well as to test the framework for soundness – that is, whether the statements and relationships between the Big Data activities and Data Governance practices implied in the framework are reasonable and logical.

Furthermore, a result of the survey was an overview of the status of commercial companies in addressing Big Data ethics. These results are diametrically opposed to the opinions expressed previously by the interviewed experts, and therefore require interpretation. The validation session brings up the opportunity to run these survey results against a panel of experts, that can help explain how commercial organizations are currently dealing with the ethical risks of using Big Data for decision making purposes.

The overview of the objectives of the validation session is shown in the table below:

Objectives	Research Question
Test governance framework for soundness property	4a
Validate the status of commercial companies in adopting Data Governance practices to address the ethical concerns raised by the use of Big Data	4b

For the validation of the research results, the Delphi technique was chosen. This method has been deemed appropriate to use when the problem at hand requires subjective judgements from a collective group (Crawford & Wright, 2016), such as for the validation of a framework (Holsapple & Joshi, 2002). Following the guidelines laid down by Crawford & Wright (2016), a heterogeneous group of experts with appropriate domain knowledge was selected to participate in the validation session. Due to the limited resources available, less than 5 experts agreed to participate, specifically two Data Governance experts and one digital ethics expert. The researcher chose to use experts working in the consulting industry, more specifically within the company within which he was conducting an internship. This choice was led by practicality reasons and takes into consideration the fact that the validation session might be affected by bias: in fact, external participants may adopt a customer-centred perspective, in contrast with the consultant perspective of the selected panel. The participants were asked in advance to be recorded: the recording of the session allows the researcher to listen back to the conversations and transcribe the most salient point, thus ensuring that the results of the validation round truly reflect what was said by the panel.

The Delphi technique consists in multiple iterations and facilitates the development of consensus among a group of experts concerning a certain topic. The feedback process allows the participants to review and rethink their initial statements (Hsu & Sandford, 2007). A semi-structured protocol was designed to guide the panel through the research results, while allowing space for any additional question they may have. The full protocol is visible in Annex H. The first round of validation consisted in a group session where the research results were presented to the panel in a semi-structured way to start the discussions. During the session, consensus among the panel participants was reached on all the points discussed. However, to guarantee that the researcher did not misinterpret their statements, a written overview of the main conclusions of the session – as interpreted by the researcher – was shared individually with each participant in a second round of validation. The participants were allowed to confirm whether the researcher had interpreted their opinions correctly, and to express any further comment regarding the topics discussed during the session. They confirmed that there were no misunderstandings in the way the researcher had understood and processed their statements, and

added a few comments they missed out in the first round that the researcher integrated in the final version of the validation results. A graphical overview of the methodology used to validate the research results is shown in the figure below.

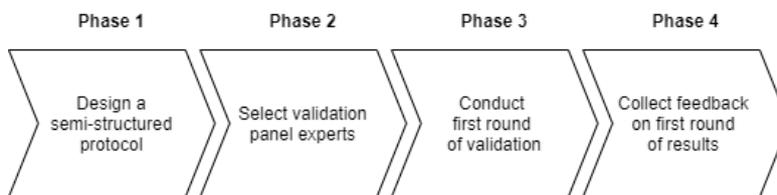


FIGURE 21 – METHODOLOGY OF VALIDATION OF RESEARCH RESULTS

4. Interview Analysis

Eight interviews were conducted with experts from relevant fields as the research topic, including ethical data compliance, information management and data ethics. The interviewees were presented with the framework resulted from the literature review and were guided through each ethical principle and relative sub-categories in order to collect feedback on the soundness of the framework itself. Furthermore, for each sub-category an example of an unethical Big Data activity was presented, in order to be validated by the experts. The interviewees were then asked to determine, among the unethical activities involving the use of Big Data for decision making, which ones they deemed to be the most risky for an organization from an ethical perspective. The answers to these questions were not only a useful input for the survey later conducted, but they also served to draft a framework of Data Governance practices for Big Data ethics.

4.1. Validation of literature framework

The data collected from the interviewees in regards to the literature review framework has been analyzed qualitatively due to the semi-structured nature of the interviews. The analysis has been conducted separately for each major ethical principle in the framework to simplify the process of transforming the load of qualitative data collected into structured results. The sub-paragraphs of this section contain the highlights of the interviews with the experts, regarding specifically their opinions and feedback on the literature framework. The complete overview of the experts opinions regarding the framework is visible in the tables of Annex I, where each row represents the opinions of a single interviewee and for each sub-category of the framework the table contains a record of whether each interviewee agreed or disagreed with it, and what additional comments or examples – if any – they had.

4.1.1. Respect for Autonomy

After the researcher introduced the ethical principle of Respect for Autonomy, along with its four categories, to the interviewees, these were asked whether they agreed or not with the definitions of the main principle and its sub categories as well as with the unethical Big Data activity related to the ethical principle. The interviewees were given the possibility to expand on the given definitions and examples and provided useful insights into the topic.

For example, for what concerns the Autonomy sub-category, the interviewees agreed with the example of the use of personalization algorithms violating this principle. For example, one interviewee stated:

“With Big Data advertising has become personalized, and it has become so smart that A: the user does not perceive that they are being influenced; and B: they cannot resist it because the Big Data system knows so much about their personality that it is able to manipulate them.”

An expert of the GDPR regulation stated that the GDPR does not actually allow the use of profiling algorithms that automate the decision making of an organization: and yet, the other interviewees did confirm that these algorithms are in place and they make choices for the user, thus violating their autonomy.

In some cases the interviewees disagreed with the researcher: one of example of this concerned the sub-category of Dignity: one interviewer stated, in response to the given example of the use of personalization algorithms to differentiate the prices presented to customers:

“No ethical principle is being violated through price differentiation based on data: in fact, using the same line of thought the differentiation of tax rates based on income would be a violation of dignity”.

While this logic is sound, one could also argue that the categorization of people can limit their ability to do certain things – which can ultimately lead to dangerous situations, Especially when a person is put into a

certain category, or shown a certain price due to factors they cannot influence, their dignity is hurt because they are being treated like the outcome of a model.

In some cases the interviewees, despite agreeing with the examples made, proposed examples from their own experience. For instance, regarding the Social Solidarity, Inclusion and Exclusion sub-category, an interviewee brought up the situation in which users' interests lead them to visit certain sites: such website will in turn only collect data and opinions about selected groups of people, leading to representativeness problems in the data.

For the principle of Respect for Autonomy, the following agreement rates were registered:

Respect for Autonomy sub-category	Agreement rate
Autonomy	100%
Dignity	86%
Informed Consent	86%
Social Solidarity, Inclusion and Exclusion	86%

Due to the high rates of agreement, the principle and its respective sub-categories were kept the same except for minor changes that were applied to it based on the interviewees' feedback. The table below shows the feedback that the interviewees had on the framework. These changes should be taken into consideration when using the literature framework in future research.

Change Log for Respect for Autonomy	Motivation
The example of the Informed Consent category is rephrased to "data collectors do not always have a lawful basis for processing personal data".	If the data collector does not have the consent of users, they could have another lawful basis which would still be a lawful way of collecting and processing personal data.
The Dignity category is rephrased to: "Human dignity means that an individual or group feels self-respect and self-worth. It is concerned with physical and psychological integrity and empowerment. (Law v. Canada (Minister of Employment and Immigration), 1999)"	The previous definition of Dignity overlapped with the definition of Discrimination.
The Social Solidarity, Inclusion and Exclusion category is rephrased to: "Some people might be excluded from the information society due to societal constructs such as cost, knowledge, social circles and class"	The previous definition of Social Solidarity, Inclusion and Exclusion was very material and did not include vaguer societal constructs. It could be for example that through Big Data applications some people might form a coalition that excludes others.

4.1.2. Nonmaleficence

The following ethical principle subject of the interview analysis is Nonmaleficence with its two sub-categories. For what concerns the sub-category of Safety, interviewees agreed that the use of Big Data may put the safety of users at risk. One interviewee stated that models are flawed and every one of them has instances in which they do not predict well; their suggestion for companies making use of these models was:

"Companies should ask what the impact of a false positive on a person is (including the potential safety risks on them), and how they are going to protect people from the flaws of the model."

The lack of a feedback loop that allows to check on the decisions generated by the analysis of Big Data is often lacking: these false positives can ultimately affect the safety of people. One interviewee proposed an example he had heard of to prove this last point: he described how some time ago in India some individuals were identified as child molesters by an algorithm and while this hadn't been problem, people went hunting for individuals and eventually beat them to death. It later turned out that these individuals were in fact not child molesters, however it was too late to go back to any decisions previously made.

For what concerns the sub-category of Discrimination and Social Sorting, one interviewee agreed with the researcher's point that the algorithms used by companies should not take discriminating characteristics into consideration, in order not to violate the ethical principle of Nonmaleficence. However, three interviewees agreed that the context is important to determine which personal characteristics are okay to use in the analysis, and which not. For example, if the purpose of the analysis will genuinely benefit the end user it should be okay to use certain characteristics; on the other hand, it is not ethical to base decisions on characteristics that people cannot influence and that will put them in a position of disadvantage. One example that was mentioned to demonstrate how discrimination might happen when using Big Data systems was about insurance companies using models that asked for the gender of people and generated a more expensive car insurance for women (which were deemed to be worse drivers by the system).

For the principle of Nonmaleficence, the following agreement rates were registered:

Nonmaleficence sub-category	Agreement rate
Safety	100%
Discrimination and Social Sorting	87%

Due to the high rates of agreement, the principle and its respective sub-categories were kept the same except for minor changes that were applied to it based on the interviewees' feedback. The table below shows the feedback that the interviewees had on the framework. These changes should be taken into consideration when using the literature framework in future research.

Change Log for Nonmaleficence	Motivation
The example of the Safety category is changed to the child molester example (described above).	The example used for the Safety category also applies to the Discrimination category, and therefore a more unique example is needed.

One interviewee pointed out that Nonmaleficence can be associated with the concept of Security: for example, if someone builds a system and he is negligent in leaving certain backdoors open, that behavior can be seen as Nonmaleficence. Furthermore, since data breaches are a real danger in Big Data, as discussed also in the GDPR, Security should be touched upon the framework. While this feedback wasn't integrated in this category, it will be discussed later on in the Data Protection category.

4.1.3. Beneficence

The next ethical principle that interviewees were presented is Beneficence with its two sub-categories. Once again, the experts were asked whether or not they agreed with the definitions and examples provided, and were able to provide further insights into this specific topic.

Regarding the ethical principle, one interviewee stated that:

"Beneficence means doing well and organizations should identify and implement possibilities to do well with data. Ethical practices should be proposed, which means that organizations should use customer data not only for their own good but also purely for the customer."

Four of the interviewed experts recognized that the current situation is one of imbalance of power, where the user gives their data away and organizations take it and do whatever it wants with it; the user doesn't see anything back from it and can only hope they won't be harmed too much by the consequences. Therefore, companies using people's data strictly for themselves in their own interest violate the ethical principle of Beneficence. The fact alone that companies own huge amounts of personal data creates an imbalance of power, which is why we could call this type of imbalance an *Imbalance of Information*, defined by a company collecting and using user data without offering any service back to the user.

One interviewee however stated that if the law works the way it should work the situation would be more balanced. He added:

“The law should have the responsibility to create a situation of [power] balance: the purpose of the GDPR for example is to give users control over their own data, which could lead to a more balanced situation.”

When discussing the sub-category of Value Sensitive Design, two interviewees agreed that when designing people put a normative stance into that design.

“The shaping of the design from the human is almost inevitable: when making something, people will, consciously or not, put some of their values into it”.

Various solutions were proposed to counterbalance this effect. One of these counterbalances for Value Sensitive Design is accountability: without accountability, one is stuck with the values of the person that designed the system in the first place; instead, if the designer can and will explain what they have done, they are open for debate and others can argue whether they agree or not with him. Furthermore, since algorithms are made by flawed humans and are subject to errors, it is important to have a human check on the algorithm and give direction to it to prevent mistakes from happening. If errors were introduced in the algorithms and people were harmed as a consequence of it, we would face a violation of the principle of Beneficence.

For the principle of Beneficence, the following agreement rates were registered:

Beneficence sub-category	Agreement rate
Imbalance of Power	75%
Value Sensitive Design	100%

As for the previous principles, due to the high rates of agreement the principle of Beneficence and its respective sub-categories were kept the same except for minor changes that were applied to it based on the interviewees’ feedback. The table below shows the feedback that the interviewees had on the framework. These changes should be taken into consideration when using the literature framework in future research.

Change Log for Beneficence	Motivation
The definition of Value Sensitive designed is changed into: “Value sensitive design means to find out what the values a product needs to fulfill for all stakeholders and consciously design to achieve certain values for the stakeholders, in a benevolent kind of way.”	The previous definition of Value Sensitive Design did not make the relation with the Beneficence principle explicit enough.
The Big Data activity for the Imbalance of Power sub-category is changed into: “The lack of transparency on what data is collected and how data is used creates an imbalance of power between the Data Collector and the Data subject.”	The activity that was previously described was not exactly an activity, but rather a definition of what Imbalance of Power is.

4.1.4. Justice

The next ethical principle analyzed by the experts is Justice and its sub-category Equality and Fairness. One expert confirmed the researcher’s point of view by stating that Justice and Fairness are basic human rights which might be violated through the misused of data. The interviewees proposed several examples of how the use of Big Data systems might lead to situations of unfairness. Two interviewees raised the case of the Dutch supermarket Albert Heijn, which sends customers personalized discounts based on information about their previous sales activity: this means that one customer might have to pay a different amount than somebody else based on the products they have bought before. A similar phenomenon occurs however also with online-displayed prices for airline tickets, hotel reservations and vacations. As another interviewee stated:

“Justice is about what society expects: if people expect that they pay the same price of a plane ticket as someone else, but actually don’t and they find out, they will feel cheated; on the other hand it would be fine if they expected it. [...] We should not fool people that think they are getting a fair price by making them pay double the price of their neighbor.”

For the principle of Beneficence, the following agreement rates were registered:

Justice sub-category	Agreement rate
Equality and Fairness	100%

As for the previous principles, due to the high rates of agreement, the principle was kept the same except for minor changes that were applied to it based on the interviewees' feedback. The table below shows the feedback that interviewees had on the framework. These changes should be taken into consideration when using the literature framework in future research.

Change Log for Justice	Motivation
In the Big Data activity of the Equality and Fairness sub-category, the wording "segmenting the population into groups" was removed and replaced with "The practice of personalization may cause different opportunities or information to be presented to the users – thus impacting the equality of people".	The wording "segmenting the population into groups" is a clear reference to the Discrimination and Social Sorting category, so the example should not allow this overlap of concepts to show.

During the discussion of the Justice principle, an expert made the point that laws and ethics are based on the geographical location and the culture of a country: they recommended to make clear that the research is based on a European idea of what is ethical and what not, and also based on European laws of privacy and data regulations. This recommendation was put in place throughout the research, but not specifically on this part of the framework.

4.1.5. Privacy and Data Protection

The last ethical principle of the framework discussed with the interviewees is Privacy and Data Protection and its eight sub-categories. One expert argued that Privacy and Data Protection is not on the same level as the other ethical principles:

"Privacy or data protection are not an end goal in themselves: by defending privacy you can have more autonomous people, there will be no abuse and society will be more fair. Aspects of this principle such as purpose specification or data minimization are different than things like equality or safety, as they are a technical mean to achieve a certain objective."

However, the other interviewees agreed on it being a category of its own, especially after both privacy and data protection became such hot topics thanks to the GDPR regulation. For each sub-category, the experts expressed their – sometimes contrasting – opinions. For example, regarding the Collection Limitation and Retention sub-category, two interviewees stated that data gathering without limitation raises privacy concerns: collecting lots of personal user data can be unethical in the sense that there can be consequences for the users if things happen to the data that the collect is not in control of (such as the loss of data). Restricting the collection of data can on the other hand mean for a company to potentially lose competitive edge. One digital law expert stated that the collection limitation is not a choice, but it is a principle demanded by the GDPR, which tries to prevent the collection of unnecessary data: the regulation states that companies can only ask to collect data on the data subject which is relevant for its purpose.

For the sub-category of Transparency, three interviewees agreed that when the algorithm is a black box, it can raise privacy – and therefore ethical – concerns: however, it is possible that even the developers that create the algorithms do not understand completely how they work, which makes it difficult to be transparent to the public and ensure that decisions can be explained. One expert argues that algorithms are going to progressively become more complex and self-thought, eventually reaching the point in which we won't be able to explain them anymore. He also stated that:

"A countermeasure for the opaqueness of algorithms is the integration of the human in the process, in a way that the human can do arbitrage and challenge the algorithm. This would also allow the integration of ethical principles in the data process."

Transparency is also mentioned in the GDPR, which states that "data must be processed lawfully, fairly and be transparent".

When discussing the sub-category of Anonymity, an interviewee told that it can be easy to trace back to a certain individual that matches certain characteristics from a data set, which makes anonymity difficult to enforce. Combining datasets together also facilitates the process of tracing back to a certain person. Furthermore, it is unethical for a company to promise his users anonymity of their data, when in fact they own enough information about them which makes identification possible. Despite this, the possibility of de-anonymizing a person is less unethical than intentionally trying to identify a person from a dataset. Nonetheless, a digital law expert referred back to the GDPR by stating that the data regulation imposes a very high threshold for data to be anonymous. Investigations have shown that even when identification chances are deemed to be very low, it is not quite the case when dealing with Big Data: this translates in higher efforts that need to be put in by companies to anonymize user data, as removing a name and an address is not sufficient.

For the principle of Privacy and Data Protection, the following agreement rates were registered:

Privacy and Data Protection sub-category	Agreement rate
Collection Limitation and Retention	50%
Data Quality	62%
Purpose Specification	50%
Use Limitation	25%
Transparency	62%
Individual Participation and Access to Data	62%
Anonymity	62%
Individual Privacy	25%

The agreement rate for this principle was on average lower than the others. While in general the interviewees went along with the sub-categories of this principle, some noted that there was no direct mapping between the GDPR principles and the categories in the framework, which led to the suggestion of replacing the sub-categories with the GDPR principles. The sub-categories were initially built on the basis of the Directive that was replaced in 2017 by the GDPR: while the researcher did establish a connection between the sub-categories and the newest GDPR regulation, the interviewees suggested that replacing the sub-categories with the GDPR principles should help avoid discussions regarding the process of mapping the Directive principles to the GDPR's. Furthermore, the GDPR is currently recognized as the common language being spoken by data experts and its principles are quite good and established, enough to the point where it is justifiable to replace the sub-categories of the Privacy and Data Protection principles with the GDPR principles.

Furthermore, a big debate was raised by the sub-category of Individual Privacy: in fact, three experts stated that Individual Privacy should not be a category of its own due to an overlap with the other Privacy and Data Protection categories. As a consequence, the researcher decided to remove this sub-category from the framework entirely.

Lastly, an interviewee raised the issue that, despite its name, data protection does not explicitly pop up in the Privacy and Data Protection principle. The importance of physically protecting data from potential breaches is not only raised by the GDPR, but also by an expert who stated that:

“Companies need to introduce measures to mitigate the harms of potential data breaches (thus reducing the chance of data breaches) but also do as much as possible by design to prevent harm being the result of a data breach.”

The table below shows a complete overview of the feedback that the interviewees had on the framework. These changes should be taken into consideration when using the literature framework in future research.

Change Log for Privacy and Data Protection	Motivation
The principle Privacy and Data Protection was renamed to “Information Privacy and Data Protection”.	Privacy consists of eight basic types including bodily privacy, intellectual privacy, etc., with informational privacy overlapping with all these types of privacy we can distinguish (Koops, et al., 2017). The principle Privacy and Data Protection focuses on Informational Privacy, in that it focuses on data that holds information. The

	other, deeper layers of privacy can be instead seen as part of the Autonomy or Dignity principles.
The definition of Privacy and Data Protection was changed to: "Privacy refers to the right of an individual to have control over the access and use of his or her personal information, as well as the right to have his or her personal information safeguarded and protected."	Privacy not only refers to the right to have control over the access of personal information, but also over its use. Also, the previous definition wasn't inclusive of the data protection part of the principle.
The sub-categories of the Privacy and Data Protection principles were changed into the principles of the GDPR.	There was no mapping between the GDPR principles and the previous categories in the framework.

4.1.7. Revised framework

Based on the feedback received by the interviewees, the literature framework was changed accordingly. While the first four principles of Autonomy, Nonmaleficence, Beneficence and Justice, along with their respective sub-categories, weren't subject to any changes, the principle of Privacy and Data Protection was radically reorganized in order for its sub-categories to match the GDPR principles. A mapping of the old sub-categories of the Privacy and Data Protection principle to the GDPR principle can support the process of reorganization of the principle. This is done on the basis of the topics covered by the old sub-categories, which can be related to the statements of the GDPR principles. The table below shows the mapping between the framework's sub-categories and the GDPR principles (which will replace the old categories in the framework).

Old sub-category	GDPR principle	Motivation
Collection Limitation and Retention	Privacy by Design, Right to be Forgotten	Privacy by Design is comprised of a data minimization article that forces data collectors to limit the collection of data to the minimum necessary. The Right to be Forgotten discusses the data controller's obligation to not retain personal data if it is not necessary for its purpose anymore.
Data Quality	Privacy by Design	The Privacy by Design controls the quality of data collected from data subjects.
Purpose Specification	Right to Access	The Right to Access describes the right of the data subjects to know for which purpose their data will be processed.
Use Limitation	Privacy by Design	Privacy by Design is comprised of a data minimization article that forces data collectors to limit the processing of data to the minimum necessary.
Transparency	Right to Access	The Right to Access controls the transparency of data processing.
Individual Participation and Access to Data	Right to Access, Right to be Forgotten	The Right to Access allows the data subject to access information regarding the purpose of the analysis and the categories of personal

		data concerned. The Right to be Forgotten allows the data subject to have their personal data erased under their request.
Anonymity	/	The GDPR does not apply to anonymized information. This may incentivize some organizations to process data anonymously (UCL, 2017).

The mapping justifies the transition to the GDPR language, as all the topics of Privacy and Data Protection included previously in the framework are covered by the GDPR principles. In addition to these, the categories of Breach Notification, Data Portability and Data Protection Officers are introduced in the framework under the principle of Informational Privacy and Data Protection. These correspond to GDPR principles which had previously not been included in the framework, but which have been described in the literature review.

The framework served the purpose of identifying risky, unethical Big Data activities on the basis of major ethical principles and their respective sub-categories. Given that during the interviews such activities were validated by the experts, despite the changes in the framework structure the following steps of the research proceeded taking such activities into account. However, future researchers wishing to use the framework to make a new assessment of unethical activities involving the use of Big Data should consider using the revised structured shown in the figure below.

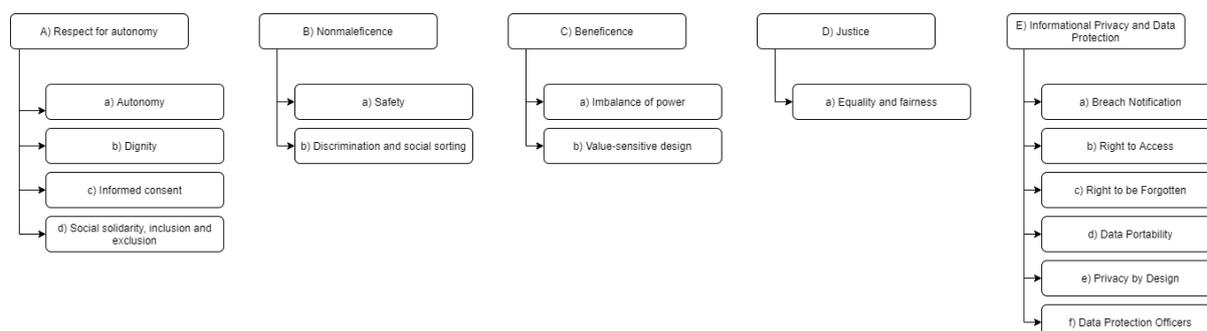


FIGURE 22 - REVISED LITERATURE FRAMEWORK STRUCTURE

4.2. Riskiness level of the Big Data lifecycle phases

The interviewees were asked to express their opinion regarding which phases of the Big Data lifecycle are the most important to address due to the ethical riskiness associated to them. For each phase, the priority level expressed by the experts was registered in a table, as seen below. The third column corresponds to the number of times the experts associated a certain priority level to a certain phase.

Phase	Choice number	Amount
Data Collection	1 st	5
Data Cleansing	2 nd	1
Data Analysis	2 nd	2
	3 rd	1
Data Visualization	2 nd	1
Decision Making	1 st	1
	2 nd	1
	4 th	1
Data Archival	3 rd	1

Data Deletion	3 rd	1
Security and Data Protection	3 rd	1

One interviewee expressed the importance of addressing the problem of Big Data ethics from a security point of view, hence the inclusion of Security and Data Protection in the list. However, being data security an overarching layer in the Big Data lifecycle it cannot be treated as a phase of its own; furthermore, being data security out of the scope of the research it was discarded from the successive calculation of the lifecycle phase priority list.

In order to determine a list of Big Data lifecycle phases ordered by priority, the scoring was calculated on the basis of the following scale:

Choice	Points
1st choice	5
2nd choice	4
3rd choice	3
4th choice	2
5th choice	1

This resulted in a newly ordered list of Big Data lifecycle phases, which are now not listed by logical succession but by the priority associated to it due to their ethical riskiness level. The overview of such list, alongside the total points scored by each phase, is displayed in the table below:

Ethical risk level	Big Data lifecycle phase	Points scored
1	Data Collection	25
2	Data Analysis	11
2	Decision Making	11
3	Data Visualisation	4
3	Data Cleansing	4
4	Data Archival	3
4	Data Deletion	3
5	Data Preparation	0
5	Data Organization	0
5	Data Integration	0

The results show that Data Collection is the lifecycle phase clearly associated with the highest ethical risk by the experts, scoring 25 points in total. The Data Analysis and Decision Making phases follow with 11 points each. The phases of Data Preparation, Data Organization and Data Integration were not mentioned a single time by the interviewees, thus they were deemed of minor importance and the researcher excluded them from further analysis in the successive stages of the study.

4.4. Data Governance standard for Big Data ethics

Together with determining which Big Data lifecycle phases are the riskiest from an ethical perspective, the interviewees were asked to propose Data Governance practices meant to address the unethical activities associated with the such risky phases.

For example, concerning the Data Collection phase, an interviewee mentioned that in order to contrast the negative ethical effect of data collectors not asking for users' consent in a specific, unambiguous way, companies should practice openness towards their customers, but also towards the authorities when it comes to showing compliance to regulations such as the GDPR. Another expert, after stating that he deemed the Data

Analysis phase one of the riskiest of the lifecycle, referring to the activity of making predictions from flawed Big Data models – which put the safety of users in danger, he proposed as a countermeasure the introduction of standardized operations within the organization to evaluate how the algorithms operate. The statements of the interviewees were gathered together to create the draft of a Data Governance standard that could be used as a tool for addressing unethical activities that might occur throughout the Big Data lifecycle within commercial organizations.

The figure below depicts a section of the framework and serves to show the structure that the framework follows. In the shown example, the unethical Big Data activity associated with the Data Collection is the act of data collectors of not ask users for consent in a specific, informed and unambiguous way. Two Data Governance practices have been identified to address this activity, namely being accountable for the data collection and being open towards customers and authorities.

Data Collection	Big Data Activities	Data Governance Practices
	Data collectors not asking users for consent in a specific, informed and unambiguous way.	Accountability for the data collection. Openness towards customers and authorities.

FIGURE 23 – STRUCTURE OF DATA GOVERNANCE FRAMEWORK FOR BIG DATA ETHICS

In some cases, a single practice was associated to an activity, and in a few exceptions some practices for the ethical handling of data were suggested without clarifying which activities these were meant to address.

Furthermore, in order to expand on the pool of Big Data activities and Data Governance practices, the paper Data Science Data Governance (Kroll, 2019) was used as a further source on information to build up on the knowledge collected through the interviews. One of the examples mentioned in the paper concerns the Data Analysis phase, where the author brings up the problem of investigating sensitive questions using company data; he proposes as a governance measure to address it the designation of a cross-functional review board responsible for examining the details of data analysis. The information present in the literature was combined to the information retrieved by means of the interviews to construct a first version of a Data Governance standard for Big Data ethics, where risky and unethical Big Data activities and corresponding Data Governance practices that addressed them are categorized by the phases of the data lifecycle. The lifecycle phases do not appear in the framework in a random order, but they are sorted by ethical risk level – as shown in section 4.2. The complete framework is visible in Annex J.

4.5. Relations between categories

When going through the framework during the interviews, the experts recognized relationships between the sub-categories of the ethical principles. While these relationships were not embedded in the Data Governance standard, the fact that they were recognized suggests that the concepts analyzed in the framework do not work in isolation, but they affect each other. Thus, violating one ethical principle might result in negative ethical consequences concerning a different principle. Conversely, living by one ethical principle might positively affect another principle.

For example, two interviewees pointed out that the act of profiling people (and discriminating them by doing so) in the wrong way can put their safety at risk. Also, when data is misused to discriminate on factors such as gender, age, etc. a violation of dignity is happening. This shows that the concepts of Discrimination, Safety and Dignity are connected to each other and one can influence another.

Two experts also recognized a relationship between the concepts of Imbalance of Power and Transparency, stating that the lack of transparency in the way data is collected and used generates the situation of power imbalance. This also means, however, that transparency can be used to correct the imbalance of power, because it allows the users to react to such situation: by exercising transparency on their data processes, organizations can give users the power to take back control of their personal data.

A complete overview of the identified relations between the sub-categories of the ethical principles is visible in Annex K. Where possible, the comments of the interviewees were included to motivate the identified relation.

5. Survey Analysis

5.1 Descriptive data

The survey ran for three weeks and collected a total of 31 responses. First, it is important to gain a more clear understanding of the population of the survey, by investigating information such as the industry they work in, their role within their companies and their general experience.

39% of the respondents work in financial services; transportation is the second most represented industry in the survey – with 16% of respondents working in it, followed by the telecommunication industry, which represents a little less than 10% of the population of respondents. All other respondents work in a multitude of industries, such as construction, IT, manufacturing and food, that all belong in the category of commercial companies. A complete overview of the industries represented in the survey is shown in Figure 1. Given the small sample sizes for each industry, it is difficult to generalize the survey results for every company belonging to each industry included in the survey: thus, the results that will follow can only say something about commercial companies in general, but not about an industry in particular.

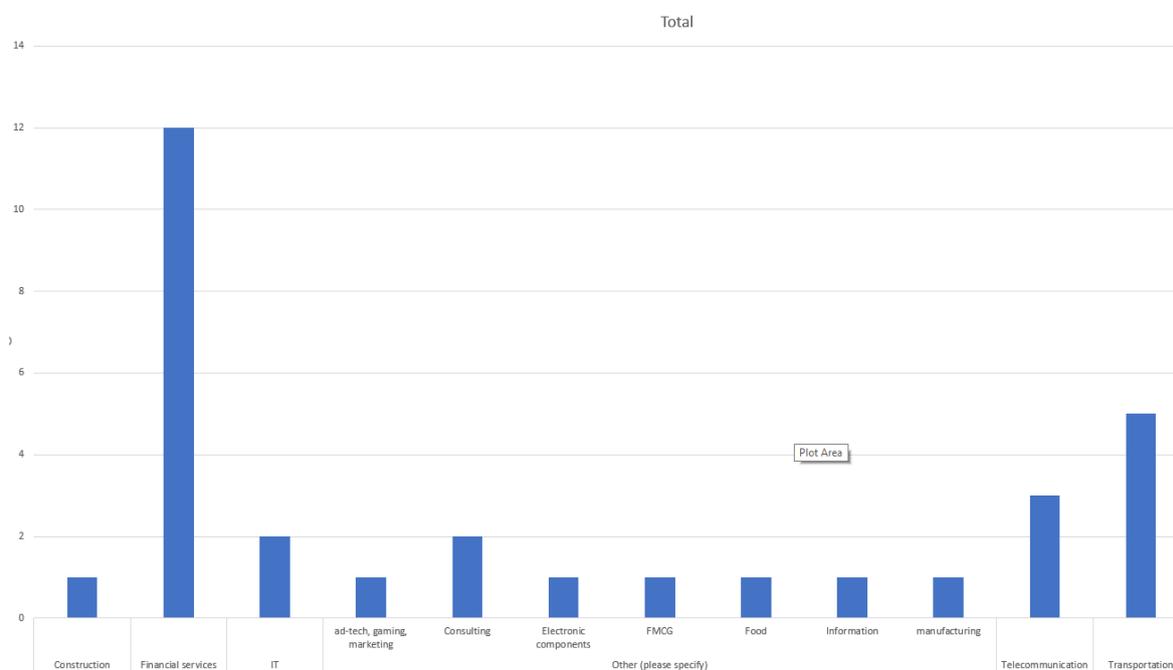


FIGURE 24 - INDUSTRY OF RESPONDENTS

As for the size of the companies the respondents work in, the survey shows to represent for the major part large enterprises comprised of more than 250 employees. Only one respondent works in a medium-sized enterprise, whereas small and micro enterprises are not represented in the survey population. This information is important to keep in mind when addressing the current status of commercial companies in addressing Big Data ethics, since the survey results are only able to draw a picture for large enterprises.

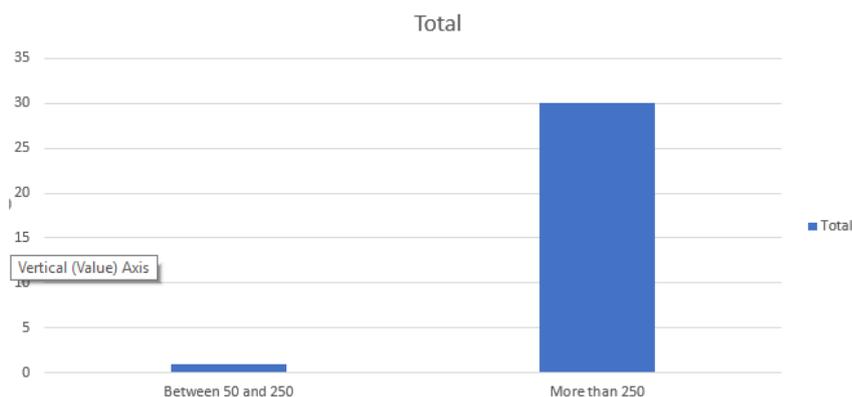


FIGURE 25 – COMPANY SIZE OF RESPONDENTS

As for the role that the respondents occupy in the companies they work in, the initial goal of the survey was to collect information from a variety of perspectives. This was successfully achieved, as shown in Figure 3. 22% of the respondents are Data Protection and Privacy Officers, 19% of the respondents are Data Scientists and Data Analysts, 13% are IT and Security architects, 10% are Technology Consultants, 6% are Data Governance Officers, 6% are Data Engineers. The remainder of professions cover a multitude of fields relevant for the research: law, business intelligence, privacy and data.

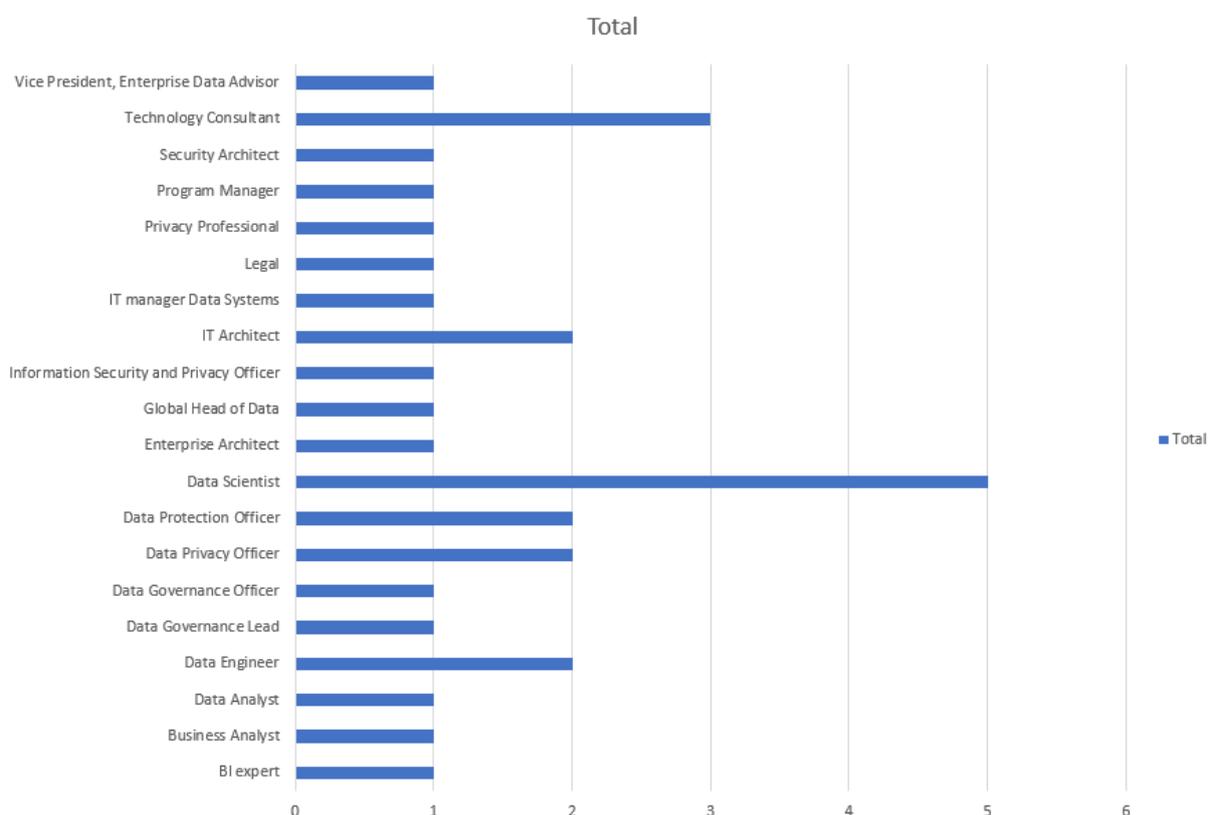


FIGURE 26 – JOB ROLE OF RESPONDENTS

Lastly, the respondents were asked to disclose their level of expertise, which is based on the years of experience they have in their respective fields. Almost 70% of the survey population has more of 5 years of experience in their field, with the remainder 30% having less than 5 years of experience. This can be considered a good balance because on one hand the respondents with many years of experience may give more elaborate and complete answers, while on the other hand younger employees may have a more fresh perspective on the topic of Big Data ethics. A detailed overview of the experience levels of the respondents is shown in the figure below.

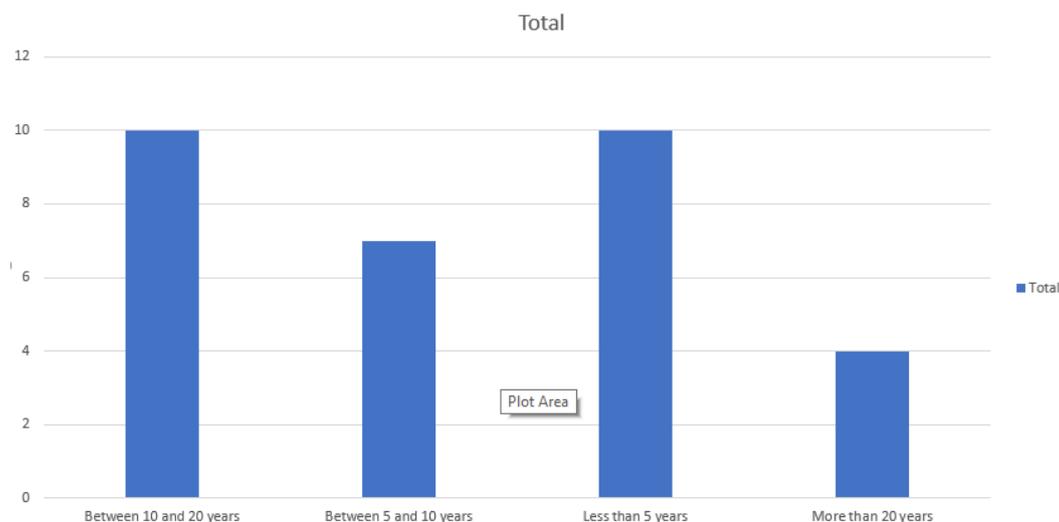


FIGURE 27 – EXPERIENCE YEARS OF RESPONDENTS

5.2 Survey results

Once more context is given about the type of respondents that participated in the survey, it is now possible to dive into the essential blocks of the questionnaire: the blocks concerning each phase of the Big Data lifecycle. The same questions were asked for each phase, and based on whether these questions were multiple choice or open-ended, the resulting data was analyzed differently.

The respondents were asked whether they have encountered ethical problems concerning a specific lifecycle phase, when using Big Data for decision making purposes, and whether the company they work for has any Data Governance practices in place to address such concerns. This information was collected through multiple choice questions, which were then analyzed in a quantitative way.

Then, by means of open-ended questions the respondents were asked to describe in detail which ethical problems, associated with a specific data lifecycle phases, they have encountered before; furthermore, after being presented with a known list of governance practices to address the ethical problems associated with a specific phase, they were given the possibility to elaborate on additional measures they would take to handle data ethically. This data was analyzed qualitatively and compared to the Big Data activities and Data Governance practices in the framework produced after the interviews. The new information collected by means of the survey was used to make changes to the framework, by adding new activities or practices to the list or by rephrasing existing ones.

Lastly, the respondents were asked contextual questions regarding the research topic: qualitative data was collected regarding their perceived risks of unethical data usage. Furthermore, the surveyed people were asked whether they place more responsibility in the law or in enterprises themselves when it comes to addressing Big Data ethics and thus enforcing ethical data usage within commercial companies.

The sub-sections down below go through the survey results concerning each lifecycle phase. Follows the analysis of the contextual questions.

5.2.1. Data collection phase

When asked if they ever encountered ethical problems associated with the Data Collection phase in their company, 48% of the respondents stated that they experienced such problems, while 52% haven't.

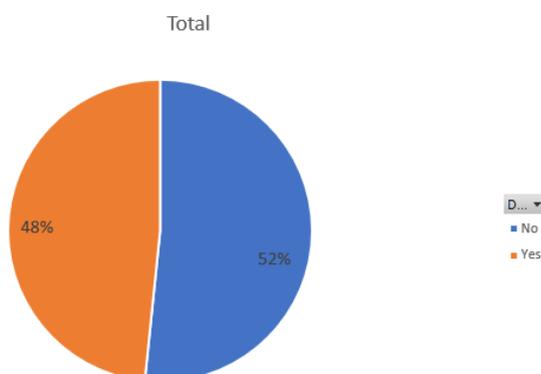


FIGURE 28 – ENCOUNTER OF ETHICAL PROBLEMS ASSOCIATED WITH DATA COLLECTION

When comparing this percentage to the percentage of respondents that have practices in place within their companies to address and deal with the ethical problems associated with the Data Collection phase, 77% of the respondents stated that they do have such practices in their company (while 16% are not sure whether they have any in place, and only 7% state that they do not have any in place at all).

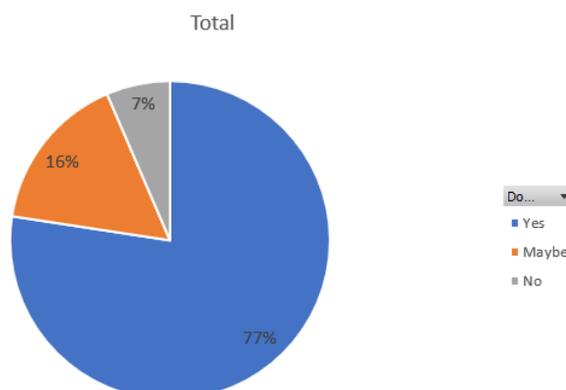


FIGURE 29 – USE OF PRACTICES FOR DATA COLLECTION PHASE

The reasons for the discrepancy within these two answers is unknown. However, we can observe that among the 16 respondents that stated they have never encountered ethical problems when dealing with the collection of data, 15 of them also stated that they have (or may have) practices in place to address such problems: this sounds logical in light of the fact that if the practices in place are as effective as they should be, no ethical problems should be experienced. On the other hand, among the 15 respondents that stated they have encountered ethical problems when dealing with the collection of data, 14 of them also stated that they have (or may have) practices in place to address the ethical problems concerning this phase: this can suggest that, despite the fact that practices to deal with these problems exist within their companies, these are not effective enough to prevent them completely from being encountered, and that perhaps additional practices are needed to counteract the effect of the ethical problems.

Changes to the framework

When it comes to the Big Data activities concerning the Data Collection phase, no additional activity was identified by the respondents: in fact, despite not giving the respondents any background information about the risky data collection activities known by the researcher, the problems identified by them match the ones previously listed in the governance framework. Some responses did serve to expand on the description of

certain activities and make such description more inclusive and complete.

For example, a respondent confirmed that ethical problems may arise by sharing data among parties: he specified that this problem is aggravated when data points are received through external suppliers but the 'opt-in' to use this data is not transparent to the organization purchasing to the data for their own use. This causes the company acquiring the data to not know what the user agreed to, with consequences for the user who might see his personal data used for purposes that he did not give his permission for. Another respondent stated that it is often unclear what data is collected and what purpose the data will be used for. The gathering of data without limitation exacerbates this problem, because it causes organizations to lose track of all the personal data they possess: this ultimately affects the user, who is also unaware of what information organizations have about them. The table below illustrates all the changes made to the Big Data activities: the light green color indicates rephrasing of existing Big Data activities, while the non-highlighted cells correspond to the activities that were not changed.

Data Collection Phase	
Old Big Data Activities	New Big Data Activities
Data collectors not asking users for consent in a specific, informed and unambiguous way.	Data collectors collecting data in illegal ways, for example by not asking users for consent in a specific, informed and unambiguous way, and by not knowing for what purpose the data itself will be used.
Excluding individuals from the data, causing the data sets to not be representative of the population and individuals not receiving the right offers.	Excluding individuals from the data, causing the data sets to not be representative of the population and individuals not receiving the right offers.
Data gathering without limitation having consequences on the user (e.g. if the data is lost)	Data gathering without limitation causes companies to not be aware of what data is being collected, with consequences on the end user.
Sharing data within and outside a company, causing the user to lose control of it.	Sharing data within and outside a company in a non-transparent way, causing the 'opt-in' to use the data to be unknown by the company acquiring the data and causing the user to lose control of its personal data.
Collecting personal data can affect customer trust and the company's reputation, and puts the company the company at risk of legal noncompliance. (Kroll, 2018)	Collecting personal identifiable information that violates people's privacy can affect customer trust - especially when uncommon types of data outside of customer expectations are collected - and the company's reputation, and puts the company the company at risk of legal noncompliance. (Kroll, 2018)

As for the Data Governance practices suggested by the respondents, while some overlap was identified with the practices in the governance framework, some additions were proposed.

For example, in order to address the problem of the illegal collection of data, two respondents suggested that organizations should apply the GDPR requirements to make sure they have the right to collect personal user data, and that the users are aware of the scope of the data collection and the methods through which this is executed. Two respondents suggested that in order to preserve customer trust, their privacy should be guaranteed by adopting anonymizing techniques. The table below shows an overview of the changes to the practices referring to the Data Collection phase: the dark green color indicates new practices that were added

to the list, the light green color indicates rephrasing of existing Data Governance practices and the non-highlighted cells correspond to practices that were not changed.

Data Collection Phase	
Old Data Governance Practices	New Data Governance Practices
Openness towards customers and authorities.	Openness towards customers and authorities: done through informing the data subject in a clear, understandable way how their data is collected and how it will create value for them.
	Application of GDPR requirements, involving data collectors making sure they have the right to collect personal user data, and that the users are aware of the scope and methods of data collection.
Reducing error of data by looking for biases in the way data is collected.	Reducing error of data by looking for biases in the way data is collected. This can be facilitated by training employees to handle data ethically and to be aware of such biases.
Defining a data strategy to take control of the data collection.	Defining a data strategy to take control of the data collection.
	Auditing the types of data collected and how it is secured.
Accompanying data with information about its provenance and processing. (Kroll, 2018)	Accompanying data with information about its provenance and processing. (Kroll, 2018)
Evaluate data for fidelity to the phenomenon under consideration. (Kroll, 2018)	Evaluate data for fidelity to the phenomenon under consideration. (Kroll, 2018)
	Set up a data sharing agreement to govern the exchange of data between parties.
Designating a review board responsible for approving or denying the collection of new data. (Kroll, 2018).	Designating a review board responsible for approving or denying the collection of new data. (Kroll, 2018).
	Anonymizing data sources, encrypting important identifiers and relinquishing or agglomerating certain data fields to guarantee the users' anonymity.

5.2.2. Data cleansing phase

When asked if they ever encountered ethical problems associated with the Data Cleansing phase in their company, 87% of the respondents stated that they have never experienced such problems, while only 13% of them have.

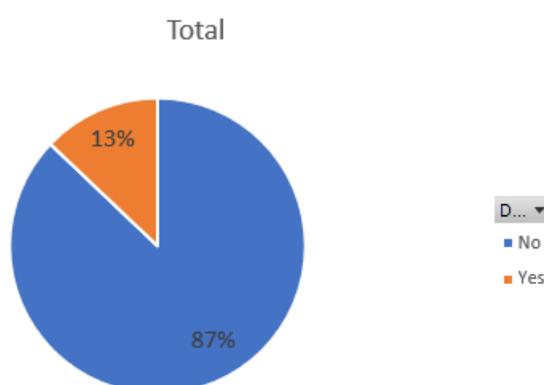


FIGURE 30 – ENCOUNTER OF ETHICAL PROBLEMS ASSOCIATED WITH DATA CLEANSING

When comparing this percentage to the percentage of respondents that have practices in place within their companies to address and deal with the ethical problems associated with the Data Cleansing phase, 42% of the respondents stated that they have such practices in place within their company, with an additional 23% who is not sure whether they do, and 35% stating they do not have such practices in place.

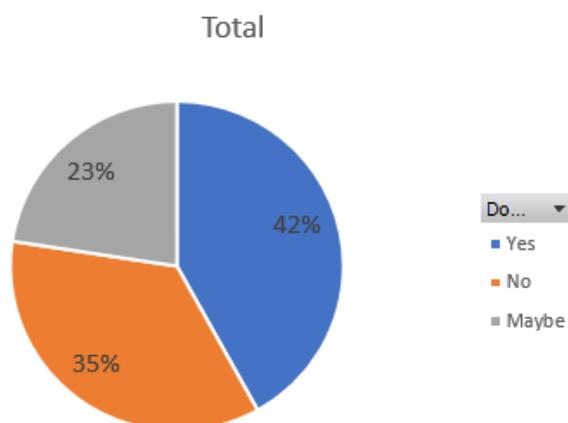


FIGURE 31 – USE OF PRACTICES FOR DATA CLEANSING PHASE

Among the 27 respondents who stated they have not encountered ethical problems in this phase, 16 of them do or may have practices in place to address them: it is justified for these 16 people to not have encountered problems, assuming the practices they have are effective in blocking the effect of unethical data cleansing activities. It is on the other hand difficult to explain why, among those who have not experienced problems before, 11 of them do not have practices in place for the data cleansing phase. This could mean that their companies are not exposed to the unethical activities concerning the data collection in the first place (perhaps because their data scientists are particularly aware of bias within the data). Alternatively, it is possible that such practices are indeed in place, however the respondents are not aware of them. For what concerns the 4 respondents who stated they have encountered ethical problems associated with the data cleansing phase, 3 of them stated they have practices in place to address them; this, once again, might be a sign of the practices not being sufficient to ensure a completely ethical cleansing of data.

Changes to the framework

When it comes to the Big Data activities concerning the Data Cleansing phase, the respondents encountered less problems than the ones identified for the previous Data Collection phase. This is in line with the limited number of activities also identified by the researcher for this phase, implying that this phase is affected by ethics less than other phases of the lifecycle.

3 respondents mentioned the risk of revealing confidential information from the cleansing of data, specifically from the operation of merging data sets together. Another respondent mentioned the risk of identifying bias while cleansing the data, bias which can have a negative impact on the end user. The table below illustrates all the changes made to the Big Data activities: the dark green color indicates new activities that were added to the framework, the light green color indicates the rephrasing of existing Big Data activities, while the non-highlighted cells correspond to the activities that were not changed.

Data Cleansing Phase	
Old Big Data Activities	New Big Data Activities
Data scientists affecting data quality with the intention to get more interesting results.	Data scientists affecting data quality with the intention to get more interesting results.
Data scientists choosing how to describe data and missing details in the world. (Kroll, 2018)	Data scientists choosing how to describe data and missing details in the world (Kroll, 2018) might introduce bias in the data.
	Associating data sets together can reveal new, sensitive and confidential data that violates the user's privacy.

As for the practices suggested by the respondents for the Data Cleansing phase, the only actionable recommendation that was given was to validate the data cleanse activities with the Data Protection Officer. The other responses were discarded because not concrete enough or not relevant for this phase. One responded however underlined the difficulty in addressing the systematic biases that humans are equipped with and might introduce in this phase: he started that due the huge difference in biases among people, there is no clear approach to address this problem. The table below shows the changes applied to the Data Cleansing part of the framework after the survey: the dark green color indicates the addition of new practices to the existing list, whereas the non-highlighted cells correspond to the practices that were not changed.

Data Cleansing Phase	
Old Data Governance Practices	New Data Governance Practices
Looking for systematic biases in the way data is cleansed and validate cleansing assumptions. (Kroll, 2018)	Looking for systematic biases in the way data is cleansed and validate cleansing assumptions. (Kroll, 2018)
Validating assumptions baked into the normalization methodology. (Kroll, 2018)	Validating assumptions baked into the normalization methodology. (Kroll, 2018)
	Align and validate the cleanse data activities with the data protection officer.

5.2.3. Data analysis phase

When asked if they ever encountered ethical problems associated with the Data Analysis phase in their company, 45% of them answers that they have not, while 55% stated that they have.

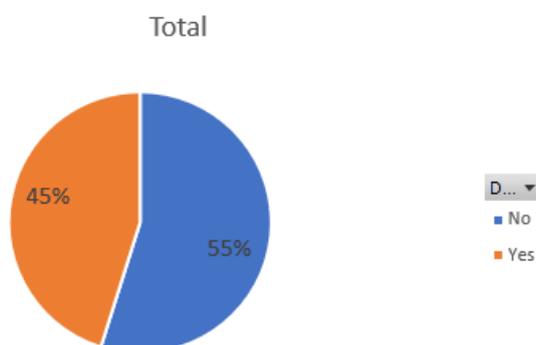


FIGURE 32 – ENCOUNTER OF ETHICAL PROBLEMS ASSOCIATED WITH DATA ANALYSIS

When comparing this percentage to the percentage of respondents that have practices in place within their companies to address and deal with the ethical problems associated with the Data Analysis phase, 42% stated that they have practices in place, with a 35% not being sure whether such practices exist within their

companies, and the remaining 23% answering that they do not have any practice in place for the Data Analysis phase.

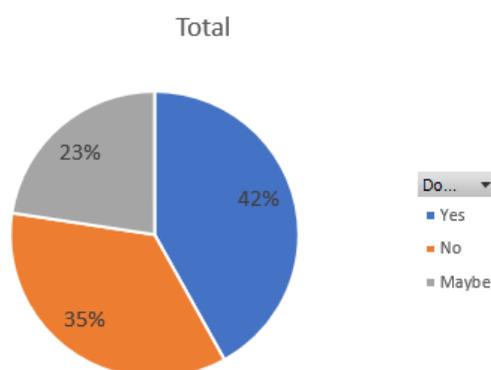


FIGURE 33 – USE OF PRACTICES FOR DATA ANALYSIS PHASE

Among the 17 respondents who have not encountered ethical problems associated with the Data Analysis phase, 14 of them do or may have practices in place to counteract these problems: this may be a sign that most of the known practices in use are effective at guaranteeing an ethical analysis of data. Among the 14 people who stated they have encountered such problems before, 7 of them also stated that they have practices in place to address them, meaning that in this case the responses would suggest an inability of the practices to address the ethical problems associated with the Data Analysis phase.

Changes to the framework

When it comes to the Big Data activities concerning the Data Analysis phase, the respondents encountered a quite substantial number of critical activities, compared to the phases previously analyzed. This is once again in line with the results of the governance framework. For the most part, the activities mentioned by the respondents corresponded to the activities already listed in the framework. The only exception was a respondent specifying that data analytics, and specifically the use of AI for the data analysis, is biased. One respondent mentioned the problem of using data for a different purpose than the one the user gave its consent for: while this problem is somewhat related to the analysis phase, consent is an issue already addressed in the data collection phase and the comment was discarded. The table below shows the rephrasing of the Big Data activities concerning the Data Analysis phase, which are highlighted in light green. The non-highlighted activities were instead not subject to any changes.

Data Analysis Phase	
Old Big Data Activities	New Big Data Activities
Use of algorithms that manipulate the user's decisions and limits their autonomy.	Use of algorithms that manipulate the user's decisions and limits their autonomy.
Predictions inferred from flawed Big Data models putting the safety of users in danger.	Errors in predictions inferred from flawed Big Data models (e.g. biased AI or analytics models, or models having outdated data as an input) putting the safety of users in danger.
Use of algorithms to target individuals in a personalized way.	Use of algorithms to target individuals in a personalized way.
Big Data algorithms that take discriminating characteristics into consideration putting users into a position of disadvantage.	Big Data algorithms that take discriminating characteristics into consideration putting users into a position of disadvantage.
Humans designing algorithms with their own perspective in mind introducing errors and biases that can harm users.	Humans designing algorithms with their own perspective in mind introducing errors and biases that can harm users.

Using black box algorithms that are difficult to understand and explain to the user.	Using black box algorithms that are difficult to understand and explain to the user.
Investigating sensitive questions using company data. (Kroll, 2018)	Investigating sensitive questions using company data. (Kroll, 2018)

As for the practices suggested by the respondents for the Data Analysis phase, many recommendations were given to address the ethical problems concerning this phase. For example it was suggested that a community of data analysis characterized by cultural, gender and ethnical diversity should be able to counteract the effects of biases throughout the Data Analysis phase. Another respondent mentioned the need of appointing a figure as a Data Protection Officer, who is responsible for enforcing measures to take to ensure a responsible and ethical analysis of data. The table below shows a full overview of the changes applied to the Data Cleansing part of the framework after the survey: the dark green color indicates the addition of new practices to the existing list, whereas the non-highlighted cells correspond to the practices that were not changed.

Data Analysis Phase	
Old Data Governance Practices	New Data Governance Practices
Standardized operations to evaluate algorithms.	Standardized operations to evaluate algorithms.
	Use a diverse analyst community (cultural, gender and ethnical diversity) to ease the removal of bias.
Establishing a common understanding of algorithms.	Establishing a common understanding of algorithms.
	Appointing a Data Protection Officer responsible for enforcing the measures needed to analyses data ethically.
Reducing error of use.	Reducing error of use.
Looking for systematic biases in the way outcomes are labelled, outliers are pruned, groupings are defined and categorical variables are encoded. (Kroll, 2018)	Looking for systematic biases in the way outcomes are labelled, outliers are pruned, groupings are defined and categorical variables are encoded. (Kroll, 2018)
	Training of professionals to be aware of the ethical and legal issues of data analysis.
	Coding reviews among data scientist to question potential ethical concerns.
Establishing effective policies and procedures to guarantee alignment between business principles and data analysis.	Establishing effective policies and procedures to guarantee alignment between business principles and data analysis.
Designating a review board responsible for examining the details of data analysis. (Kroll, 2018)	Designating a review board responsible for examining the details of data analysis. (Kroll, 2018)

5.2.4. Data visualization phase

When asked if they ever encountered ethical problems associated with Data Visualization in their company,

81% of the respondents answered that they have not experienced any ethical problems concerning this phase, with only 19% claiming they have.

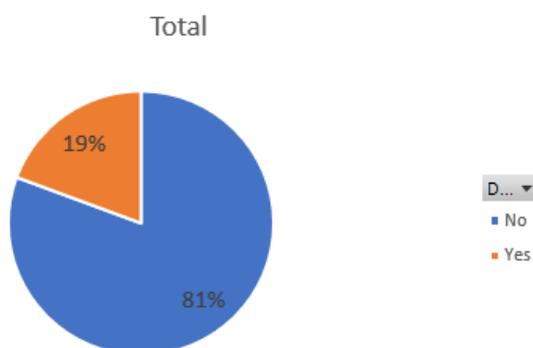


FIGURE 34 – ENCOUNTER OF ETHICAL PROBLEMS ASSOCIATED WITH DATA VISUALIZATION

When comparing this percentage to the percentage of respondents that have practices in place within their companies to address and deal with the ethical problems associated with the Data Visualization phase, 62% of the respondents said that they do not have such practices in place, with a 32% saying they do.

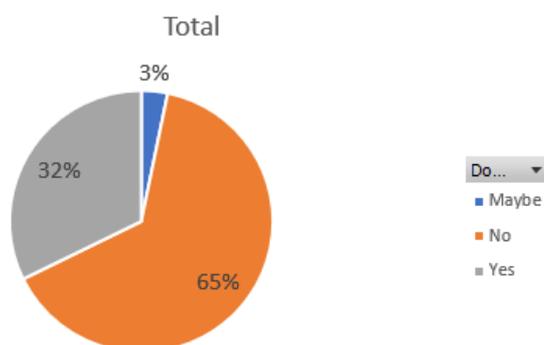


FIGURE 35 – USE OF PRACTICES FOR DATA VISUALIZATION PHASE

Among the 21 people who claimed they haven't experienced problems concerning this phase, only 8 of them stated they have practices in place. This could be explained by the fact that, even though the companies have practices in place, the respondents were not aware of them. It could also be a sign of those visualizing data being particularly aware of the biases they might introduce in doing so, and thus preventing them from showing. On the other hand, among the 10 people who claimed to have experienced data visualization problems, 7 of them said they do not have practices in place. Keeping in mind that the majority of respondents did not experience problems concerning this phase at all, these numbers might be a sign that data visualization problems are not as impactful on the data lifecycle from an ethical perspective, and therefore the impact of practices to address potential issues is low compared to other phases. This result would be in line with the fact that only one risky activity concerning this phase has been identified by the researcher.

Changes to the framework

For what concerns the Big Data activities related to visualizing data, the respondents showed to not experience many problems concerning this phase. Three respondents described an unethical activity that matches the one already listed in the governance framework: they find that when visualizing data, people tend to focus on a particular, favorable outcome instead of presenting the data set strengths and weaknesses, the assumptions hidden in the model and its accuracy. Three other respondents raised a new problem: the exposure of sensitive data concerning customers to the employees working on the data visualization. The table

below shows the overview of the changes applied to the Big Data activities for the Data visualization phase: the boxes in dark green show additions of activities to the existing list; the activities highlighted in light green have been rephrased and the non-highlighted boxes correspond to the activities that haven't been changed.

Data Visualization Phase	
Old Big Data Activities	New Big Data Activities
Humans interpreting analysis results introducing errors and interpreting the results based on their personal values.	Humans interpreting analysis results introducing errors and interpreting the results based on their personal values (bias) or in a way that favors a particular outcome - ignoring to present the model assumptions and accuracy.
	Exposing the visualized, sensitive customer data related to users to employees (especially to those who should not see it) and showing it in reports.

As for the practices suggested by the respondents for this phase, two respondents suggested that in order to address the problem of exposing employees to sensitive data, the visualized information should be anonymized; it should also be made clear by means of guidelines which employees in the company are authorized to work with customer-specific data and visualize it. Furthermore, a respondent pointed out the need to set limits on what information is necessary to know in order to be able to make a decision: this is important to do throughout the entire data process, but especially in this step since it affects the way decision making will be executed later on. While this is an important suggestion to keep in mind, it wasn't directly translated into a practice because the relation between such practice and the corresponding activity that it means to tackle is not clear enough. The table below shows the changes applied to the governance practices related to the Data Visualization phase: the practices in dark green have been newly added to the model, whereas the boxes in light green correspond to practices which have been rephrased.

Data Visualization Phase	
Old Data Governance Practices	New Data Governance Practices
Being aware of human bias which affects the way results are interpreted.	Training employees on how to interpret data and to be aware of the human bias which affects the way results are interpreted.
	Anonymizing data before it is presented to employees
	Drafting guidelines that regulate which data can be shown to which employees in the company and which practitioners are mandated to work with customer-specific data.

5.2.5. Decision making phase

When asked if they ever encountered ethical problems associated with Decision Making phase in their

company, 68% of the respondents answered that they have not experienced any ethical problems concerning this phase, with only 32% of them claiming they have.

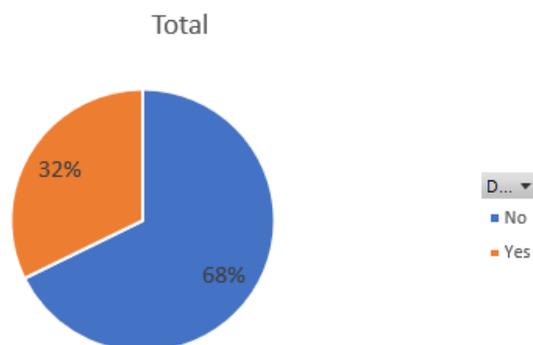


FIGURE 36 – ENCOUNTER OF ETHICAL PROBLEMS ASSOCIATED WITH DECISION MAKING

When comparing this percentage to the percentage of respondents that have practices in place within their companies to address and deal with the ethical problems associated with the Decision Making phase, 52% of the respondents said that they have such practices in place, with a 22% saying they don't and the remaining 26% not being sure.

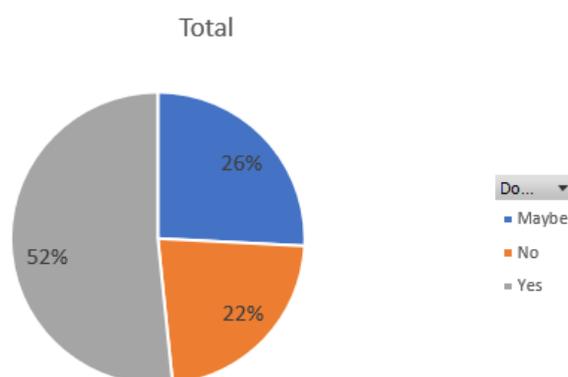


FIGURE 37 – USE OF PRACTICES FOR THE DECISION MAKING PHASE

Among the 21 people who stated they haven't experienced any problems with the Decision Making phase, 14 of them also claimed they may or do have practices in place to prevent them. The 7 respondents who claimed they do not might not be aware of them, or perhaps the ethicality of this phase is guaranteed informally by decision makers being particularly aware of making decisions that do not have negative consequences on the end users. On the other hand, among the 10 people who stated they have experienced problems with the Decision Making phase of the lifecycle, all of them stated they do or may have practices in place. The justification for this result can only be hypothesized: one possibility could be that these practices are not effective enough in preventing ethical problems from showing up; alternatively, it might be possible that the effect of unethical practices in previous phases of the lifecycle might have a negative impact on the Decision Making phase, which cancels the effect of the ethical practices.

Changes to the framework

For what concerns the Big Data activities related to the Decision Making phase, the respondents' answers were a match with the activities already listed in the governance framework. The insights of respondents were however useful for making the description of the activities more complete. For example, two respondents mentioned that two specific data-driven systems, namely AI-based systems and Machine Learning algorithms, may impact the decisions made by producing wrong decisions (e.g. automated recruiting systems may not always select the right candidates). Another problem raised by a respondent was that there is an ethical dilemma behind allowing service-oriented decisions to decide for someone else, namely the end users. This is

related to the idea that the use of Big Data analytics restricts the autonomy of individuals, who are subject to the decisions produced by Big Data systems without having to chance to intervene in the process. The table below shows a complete overview of the changes applied to the Big Data activities for the Data visualization phase: the activities highlighted in light green have been rephrased and the non-highlighted boxes correspond to the activities that haven't been changed.

Decision Making Phase	
Old Big Data Activities	New Big Data Activities
Sending individuals personalized offers and information, creating a situation of unfairness.	Sending individuals personalized offers and information, creating a situation of unfairness.
Drawing conclusion from poor quality data, with repercussions on the data subjects.	Drawing conclusion from poor quality data or externally acquired data of unknown accuracy and reliability, with repercussions on the data subjects.
Users being unaware of how predictive information is inferred and impacts them.	Users being unaware of how predictive information is inferred and impacts them.
Generating mistaken decisions and false positives, causing the putatively high-risk individuals to be treated unfairly. (Kroll, 2018)	Generating mistaken decisions and false positives on the basis of discriminating characteristics, causing the putatively high-risk individuals to be treated unfairly. (Kroll, 2018)
Using data-driven systems, which can be affected by modeling errors and whose fidelity changes over time. (Kroll, 2018)	Using data-driven systems (e.g. AI-based systems and ML algorithms), which can be affected by modeling errors and whose fidelity changes over time. (Kroll, 2018)
Deployment of insights from analysis of sensitive questions. (Kroll, 2018)	Deployment of insights from analysis of sensitive questions, leading to decisions which may contain an ethical problem or dilemma.

Similarly to the activities related to the Decision Making phase, also for the corresponding governance practices the respondents had no new insights, but with their answers they helped rephrasing existing practices in the framework list for increased clarity. For example, one respondent clarified that the internal role who should be designated to be responsible for owning the outcomes of the data analysis is the Data Protection Officer. Another respondent highlighted the importance of having a group of people, instead of a single individual, to solve ethical problems related to the Decision Making phase: the greater group is often perceived as more knowledgeable to solve problems. The table below shows the changes applied to the Data Governance practices for this phase: the practices in light green were rephrased, whereas the practices in the non-highlighted boxes were not changed.

Decision Making Phase	
Old Data Governance Practices	New Data Governance Practices
Disclosing analysis methods to guarantee the transparency. (Kroll, 2018)	Disclosing analysis methods to guarantee the transparency. (Kroll, 2018)
Monitoring the performance of a system after launch by means of black box testing to test it against unfairness. (Kroll, 2018)	Monitoring the performance of a system after launch by means of black box testing to test it against unfairness. (Kroll, 2018)
Designating an internal role responsible for owning the outcomes of analysis. (Kroll, 2018)	Designating the internal role of the Data Protection Officer, responsible for owning the outcomes of analysis.
Establishing a common understanding of how specific decisions are made (transparency).	Establishing a common understanding of how specific decisions are made (transparency).
Considering the possibility that mistaken decisions might disproportionately harm individuals or protected groups and testing for feedback loops. (Kroll, 2018)	Considering the possibility that mistaken decisions might disproportionately harm individuals or protected groups and testing for feedback loops. (Kroll, 2018)

Having data scientists validating predictions and monitoring the performance of systems after launch. (Kroll, 2018)	Having data scientists manually check the results from automated analysis by validating predictions and monitor the performance of systems after launch.
Designating a review board responsible for approving or denying the use of analytics insights. (Kroll, 2018)	Designating a review board able to gain a common understanding of ethical problems, responsible for approving or denying the use of analytics insights.

5.2.6. Decision archival and deletion phases

When asked if they ever encountered ethical problems associated with the Data Archival and Deletion phases in their company, 68% of the respondents answered that they have not experienced any ethical problems concerning this phase, with the remaining 32% claiming they have.

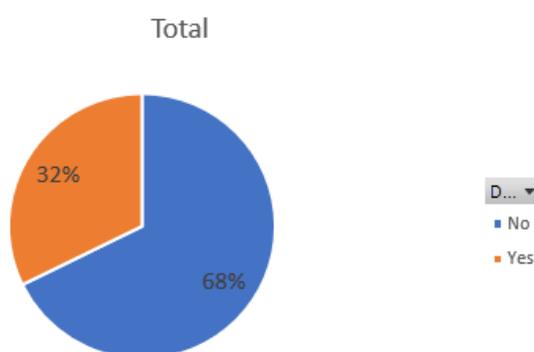


FIGURE 38 – ENCOUNTER OF ETHICAL PROBLEMS ASSOCIATED WITH DATA ARCHIVAL AND DATA DELETION

When comparing this percentage to the percentage of respondents that have practices in place within their companies to address and deal with the ethical problems associated with the Data Archival and Deletion phases, 58% of the respondents said that they have such practices in place, with a 29% saying they don't and the remaining 13% not being sure.

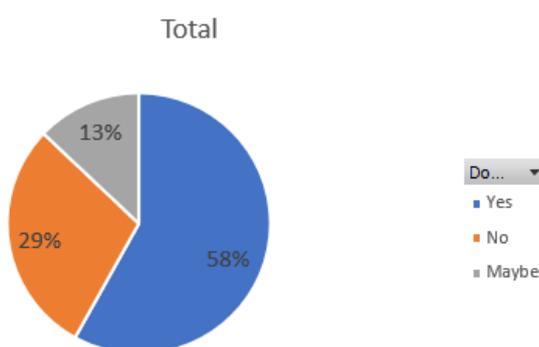


FIGURE 39 – USE OF PRACTICES FOR DATA ARCHIVAL AND DATA DELETION PHASES

When comparing these graphs together, among the 21 respondents who stated they have never experienced ethical problems concerning these phases, 12 of them also said they have practices in place to address them. In order to explain why the 9 remaining respondents said they don't have practices in place, we can hypothesize that their data process prior these stages is ethical enough that no problems are experienced when data is achieved or stored. As for the 10 respondents who answered that they have experienced ethical problems connected to the archival and deletion of data, all 10 stated that they do or may have practices in place for this phase. Since once again it is difficult to come up with an exact explanation for this phenomenon, we can only hypothesize that in these 10 cases the data processes prior to the archival and deletion of data

have let ethical problems reach these phases, and that any practice in place concerning these phase is not able to solve them.

Changes to the framework

For what concerns the Big Data activities related to archiving and deleting data, three respondents mentioned a novel problem that hadn't been considered in the governance framework: the retention of personal data past its due date with no clear business use. Four respondents mentioned the problem of deleting data: even when requested by individual users on the basis of GDPR, this process does not guarantee that data is destroyed for good. Furthermore, companies may incur in some conflict of interest when deciding whether to destroy data or not: for example, when testing for model replication the input data of the model must be retrievable at the moment of the model review, meaning that such input data is never completely destroyed. Despite these respondents raising the problem of data destruction, they were unable to propose a concrete practice as a solution to address it. The table below shows the overview of the changes applied to the Big Data activities for the Data Archival and Deletion phases: the boxes in dark green show additions of activities to the existing list, while the non-highlighted boxes correspond to the activities that haven't been changed.

Data Archival and Deletion Phases	
Old Big Data Activities	New Big Data Activities
Reidentification of user data despite promising users anonymity of their personal information.	Reidentification of user data despite promising users anonymity of their personal information.
	Storing personal data past its due date with no clear business use.

As for the practices suggested by the respondents for these two phases, three respondents mentioned the importance of establishing a clear data retention policy, which involves the act of balancing corporate, public and individual interests. Furthermore, in order for the retention process to have no loose ends, it is crucial to ensure both that data is profitably discarded and properly anonymized. The table below shows the changes applied to the governance practices related to the Data Visualization phase: the boxes in light green correspond to practices which have been rephrased, while the non-highlighted practices have not been changed.

Data Archival and Deletion Phases	
Old Data Governance Practices	New Data Governance Practices
Considering the risk that retained data could be reidentified, which depends on the type of data in question and the context in which it is being used. (Kroll, 2018)	Considering the risk that retained data could be reidentified, which depends on the type of data in question and the context in which it is being used. (Kroll, 2018)
Understanding how and why data must be retained and how it will be used, to know how it can be profitably discarded or properly anonymized to minimize risk. (Kroll, 2018)	Establishing a clear data retention policy that clarifies how and why data must be retained and how it will be used, to know how it can be profitably discarded and properly anonymized to minimize risk and make sure the retain process has no loose ends.

5.2.7 Risks of unethical Big Data usage

When asked how risky for their company is to leave ethical problems deriving from the using Big Data unaddressed, all of the respondents answered that they find it risky, with none of them answering they do not find it risky at all. The perceived level of riskiness varied among the respondents, with the majority of them finding the use of Big Data ethics for decision making purposes very risky, and 11 of them rating the riskiness on the highest level. Only one respondent selected the option 'slightly risky'. This result would imply that among all the people surveyed, all of them are aware of the risks of making decisions on the basis of analysis of Big Data, with 84% of them believing this risk is high.

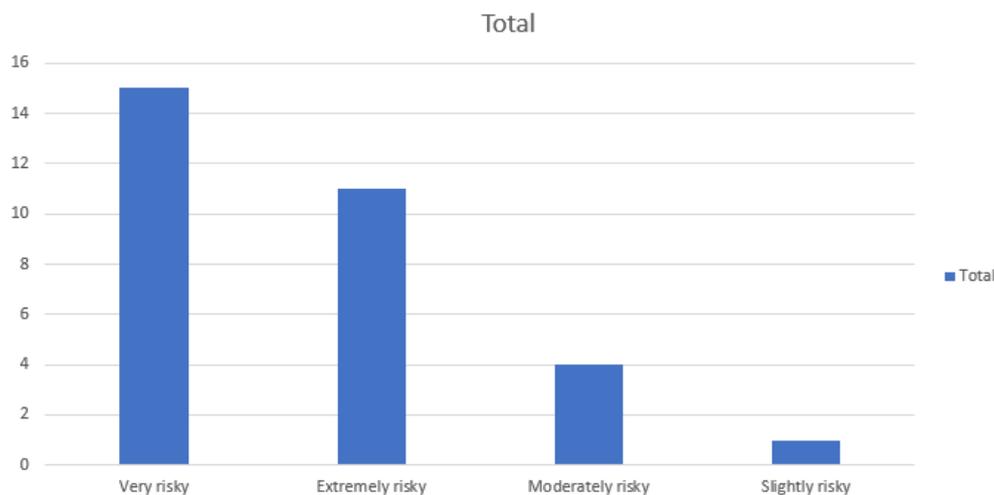


FIGURE 40 – PERCEIVED RISKINESS OF USING BIG DATA FOR DECISION MAKING PURPOSES

The respondents were also asked to explain what the risks of handling data unethically are for their respective companies. The majority of respondents (51%) stated that the major risk their company faces is reputation damage: for example, banks risk this type of damage when they are not transparent in explaining the reasoning for prediction outcomes. Using personal user data unethically could result in court cases that might damage the reputation of a company. The second biggest risk, as perceived by the respondents, is financial damage, which is strictly related to the third biggest risk: non-compliance. In fact, not being legally compliant with data protection laws, GDPR in particular, can result in penalties and regulatory fines. Financial damage can also be result of the reputation damage to the company's image. 5 of the respondents then listed various ethical risks as one of the consequences of unethical data usage: among these risks is the introduction of biases in areas such as the financial sector (e.g. giving low cost loans to wealthy clients and giving high cost loans to poorer clients, leading to an exacerbation of wealth differences among the population), which can in turn have a negative effect on individuals. Furthermore, the respondents perceived the loss of customers as another risk they might potentially face as a consequence of handling their data incorrectly: unethical handling can breach customer trust and trigger churn, particularly in competitive environments – with the ultimate effect of financial damage for the company. A complete overview of the risks mentioned by the respondents, with the corresponding frequency of mentions, is shown in the table below.

Risk	Frequency of mentions
Reputation damage	16/31
Financial damage	6/31
Ethical risks	5/31
Non-compliance	4/31
Loss of customers	4/31

What we can conclude from the responses is that there is not one individual risk that can be associated with the unethical usage of data, but that one type of damage can trigger others in a cascade effect.

5.2.8 The responsibility of the law vs companies

One of the issues addressed throughout this study is the contrast between laws and regulations on one hand, and companies on the other, in enforcing ethical behavior when it comes to using Big Data for making decisions. In the survey, the respondents were asked their opinions regarding this debate: specifically, they were asked the level of responsibility that respectively the law and companies have in establishing ethical behavior within an organization.

The majority of respondents (87%) attributed a medium to high level of responsibility to laws and regulations in establishing ethical data behavior. Only 1 of the respondents stated that they believe the law has no power in enforcing such behavior. The complete overview of responses is shown in the figure below.

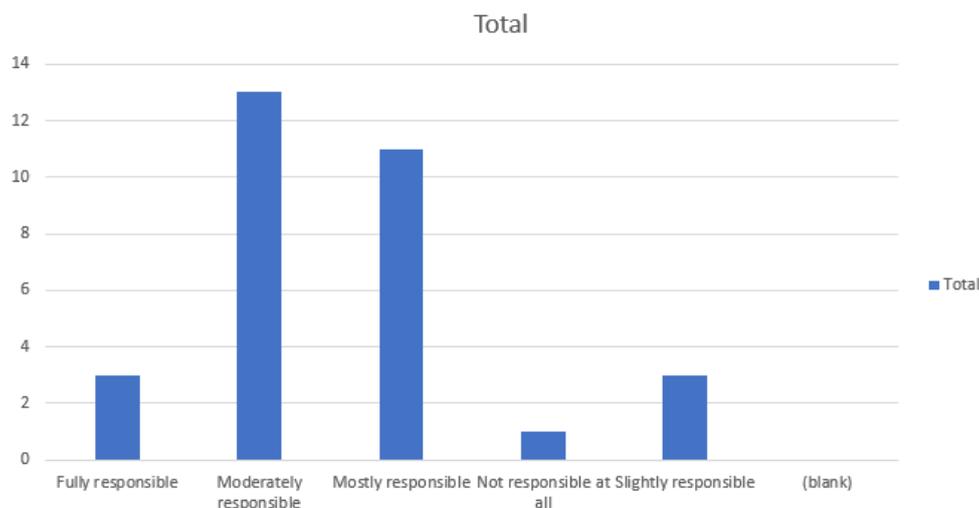


FIGURE 41 – PERCEIVED RESPONSIBILITY OF LAWS AND REGULATIONS IN ESTABLISHING ETHICAL DATA BEHAVIOR

On the other hand, all the respondents attributed some level of responsibility to organizations themselves in ensuring that data is used ethically within them. In fact, none of them stated that organizations have no responsibility at all, and 97% of the respondents attributed to companies a medium to high level of responsibility, as shown in the figure below.

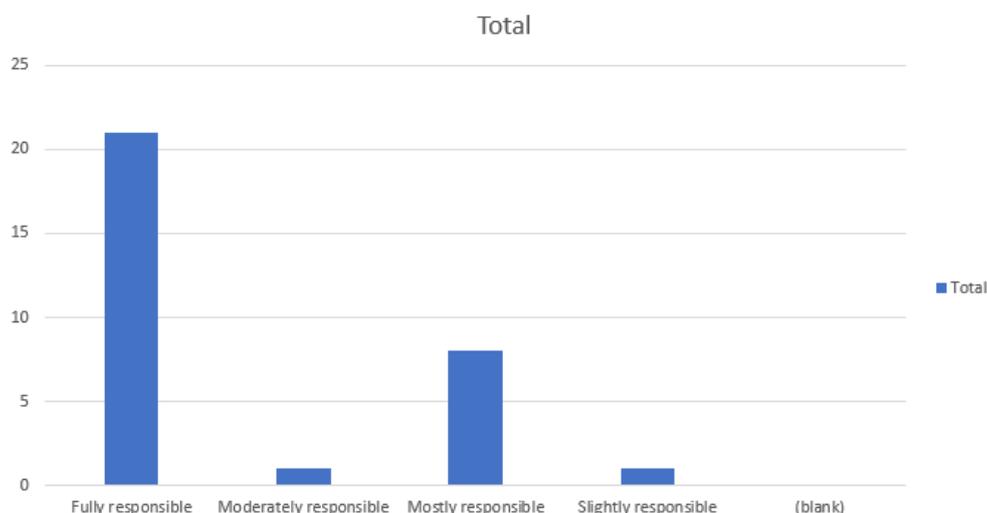


FIGURE 42 - PERCEIVED RESPONSIBILITY OF COMPANIES IN ESTABLISHING ETHICAL DATA BEHAVIOR

5.3. Status of commercial organizations in addressing Big Data ethics

One of the purposes of the questionnaire was to make an assessment of the ethical usage of data within the surveyed commercial organizations. In order to determine how well prepared these enterprises are in addressing the ethics of Big Data, the respondents were asked to rate the lifecycle phases from the phase of highest importance, to the phase of lowest importance, as perceived by them. The resulting ranking was then compared to the ranking of the lifecycle phases which resulted from the interviews with experts: this ranking corresponds to the phases that were perceived by the experts the most important to address due to their riskiness level. The comparison table below shows that there is not a match between the two rankings: both the experts and the survey respondents deemed the data collection phase more important to address than the

data analysis phase, and both deemed the cleansing phase more important than the data visualization phase, which is also more important than archiving and deleting data.

Perceived importance of lifecycle phases	Riskiness level of lifecycle phases
Data cleansing	Data collection
Data visualization	Data analysis
Data archival and data deletion	Decision making
Data collection	Data cleansing
Data analysis	Data visualization
Decision making	Data archival and data deletion

While there is a similarity of thought among the two groups, the overall difference in result could be explained by considering that the ranking of the lifecycle phases might be correct from a theoretical perspective, practitioners may feel differently.

The respondents were also asked to rank the Big Data lifecycle phases from the lowest priority phase to the highest priority phase within their companies specifically. Also in this case there was a discrepancy between the priority level of the phases within the surveyed organizations and the desirable priority ranking, as described by the interviewed experts. One significant result was, for example, the fact that while the data collection phase should be the highest priority within organizations, the survey respondents placed it on the 5th place in their ranking. Similarly, the data analysis phase, which was ranked second in terms of riskiness by the experts, is only at the 4th place in the priority list of the surveyed companies. The complete comparison between the two rankings is shown in the table below.

Level of priority of lifecycle phases within organizations	Riskiness level of lifecycle phases
Data visualization	Data collection
Data archival and data deletion	Data analysis
Data cleansing	Decision making
Data analysis	Data cleansing
Data collection	Data visualization
Decision making	Data archival and data deletion

The mismatch between riskiness level and priority level might be a sign that the companies that participated in the survey are not addressing the ethics of Big Data correctly, by putting their priorities on the wrong activities of the data process.

5.3.1. Governance practices in use within commercial organizations

Throughout the survey, the respondents were given the chance to select among the list of Data Governance practices of the framework the ones that they have in place within their companies. By deriving from these answers a list of the most commonly used practices within commercial organizations, it should be possible to make an assessment of where they stand in terms of Big Data ethics. In the sections below the results of such analysis are presented for each phase of the lifecycle.

Data Collection

For what concerns the Data Collection phase, among the nine practices that were presented to the respondents, only four of them got mentioned by a significant amount of respondents. Specifically, the practice that resulted to be the most commonly used among the surveyed companies is *Accountability for the data collection* (which was selected by 61% of the respondents) followed by *Define a data strategy* (selected by 48% of the respondents) and *Openness towards the customers and authorities* (practice used by 45% of the respondents). *Establish effective policies and procedures to guarantee alignment between business principles and the collection of data* came into fourth place with a 35% popularity. All the other practices were more or

less unpopular among the companies part of the study and they can be viewed in the table below. The exact reason why these specific practices are the most used ones is unknown and should be investigated further with a panel of experts.

Data Governance practices for Data Collection	Frequency of mentions
Accountability for the data collection	19
Define a data strategy	15
Openness towards the customers and authorities	14
Establish effective policies and procedures to guarantee alignment between business principles and the collection of data	11
Reduce error of data by looking for biases in the way data is collected	3
Design a review board with the power to approve or deny the collection of new data	3
Accompany data with information about its provenance and processing	2
Evaluate data for fidelity to the phenomenon under consideration	1
Anonymizing data sources, encrypting important identifiers and relinquishing or agglomerating certain data fields to guarantee the users' anonymity.	1
Setting up a data sharing agreement to govern the exchange of data between parties.	1

Data Cleansing

For what concerns the Data Cleansing phase, the practice that was selected the most was *Establish effective policies and procedures to guarantee alignment between business principles and the data cleansing phase*, with 42% of the respondents stating they have this practice in place; this practice was followed by *Define a data strategy to clarify how you work with data*, which was selected by 35% of the respondents. The remaining three practices got less than 5 mentions, with the practice *Validate assumptions baked into the normalization methodology* receiving no mentions at all. The full overview of the frequency of mentions for each practice is shown in the table before. As seen for the Data Collection phase, a large amount of respondents stated they defined a data strategy that makes clear how the organization both collects and cleanses data.

A validation with experts should aim to clarify why certain practices are more used than others in this phase.

Data Governance practices for Data Cleansing	Frequency of mentions
Establish effective policies and procedures to guarantee alignment between business principles and the data cleansing phase	13
Define a data strategy to clarify how the company cleanses data.	11
Look for systematic biases in the way data is cleansed and validate data cleansing assumptions	5
Aligning and validating the data cleanse activities with the data protection officer.	1
Validate assumptions baked into the normalization methodology	0

Data Analysis

Moving onto the Data Analysis phase, only three practices were mentioned by circa 30% of respondents or more, with *Accountability of data usage* being the most frequently used practices (47% of respondents stated they have it in place within their companies). *Establish effective policies and procedures to guarantee alignment between business principles and the data analysis phase* follows with 32% of preferences expressed and *Define a data strategy to clarify how the company works with data* mentioned by 29% of the respondents. All other practices were mentioned by 5 respondents or less, but they all showed to be in place within the surveyed company to a certain extent. Please refer to the table below for the complete list of practices and respective frequency of mentions. Similarly to what observed for the Data Collection phase, accountability emerged as a frequently mentioned concept for ensuring ethical handling of data.

A validation with experts is needed to clarify why certain practices are more used than others in this phase.

Data Governance practices for Data Analysis	Frequency of mentions
Accountability of data usage	15
Establish effective policies and procedures to guarantee alignment between business principles and the data analysis phase	10
Define a data strategy to clarify how the company works with data	9
Designate a cross-functional review board with the power to approve or deny the deployment of insights from the analysis of sensitive questions	5
Reduce error of analysis	3
Look for systematic biases in the data analysis phase	3
Introduce standardized operations to evaluate algorithms	2
Establishing a common understanding of algorithms	2
Training of professionals to be aware of the ethical and legal issues of data analysis.	1
Coding reviews among data scientist to question potential ethical concerns.	1
Appointing a Data Protection Officer responsible for enforcing the measures needed to analyze data ethically.	1

Data Visualization

For what concerns the Data Visualization phase, given the high number of respondents who claimed to not have any practice in place for this phase, a limited number of mentions was collected in this part of the survey. 6 respondents claimed to have the practice *Awareness of human bias which can affect the way results are interpreted* in use within their organizations, with the remainder two practices collecting one mention each, as shown in the table below. A validation with experts should aim to clarify the lack of governance practices within the surveyed organizations for the Data Visualization phase.

Data Governance practices for Data Visualization	Frequency of mentions
Awareness of human bias which can affect the way results are interpreted	6
Anonymizing data before it is presented to employees	1
Drafting guidelines that regulate which data can be shown to which employees in the company and which practitioners are mandated to work with customer-specific data.	1

Decision Making

Regarding the Decision Making phase, the practice *Establish effective policies and procedures to guarantee alignment between business principles and decision making phase* was the most popular choice, with 42% of respondents selecting it. Similarly to what seen in all previous phases, the respondents claimed to have policies and procedures in place to guarantee the alignment between their business principles and the phases of the Big Data lifecycle: this practice was in fact a popular choice for all the phases so-far analyzed. For what concerns the Decision Making phase in particular, the practices *Establish a common understanding of how specific decisions are made*, *Have data scientists validate predictions continually* and *Designate an internal role responsible for owning the outcomes of analysis* were all fairly popular, with a frequency of mentions between 22 and 29%. The table below contains more detailed information about all the practices mentioned by the respondents.

Also for this phase, the reasons why certain practices are in use more than other should be investigated further.

Data Governance practices for Decision Making	Frequency of mentions
Establish effective policies and procedures to guarantee alignment between business principles and decision making phase	13
Establish a common understanding of how specific decisions are made	9
Have data scientists validate predictions continually	8

Designate an internal role responsible for owning the outcomes of analysis	7
Disclose analysis methods to guarantee the transparency	6
Test for feedback loops and the possibility that mistaken decisions might disproportionately harm individuals	5
Monitor the performance of a system after launch by means of black box	1
Disclosing analysis methods to guarantee the transparency	0

Data Archival and Data Deletion

For what concerns the Data Archival and Data Deletion phases, 51% of the respondents claimed that they address the ethical problems concerning these phases by *Establishing a clear data retention policy*. 29% of respondents stated that they have the practice *Consider the risk that retained data could be reidentified* in place. Since all the mentioned practices are used to some extent, there is no need to investigate the usage of practices concerning this phase any further.

Data Governance practices for Data Archival and Data Deletion	Frequency of mentions
Establishing a clear data retention policy that clarifies how and why data must be retained and how it will be used, to know how it can be profitably discarded and properly anonymized to minimize risk and make sure the retain process has no loose ends	16
Consider the risk that retained data could be reidentified	9

6. Results Validation

By means of a survey, a group of data practitioners occupying a diversity of data-related roles within commercial organizations were given the opportunity to explore the Big Data activities and Data Governance practices they are exposed to; this allowed an implicit validation of the known activities and practices that constitute the Data Governance standard for Big Data ethics produced throughout this study. What emerged from the processing of the survey results is that the majority of activities and practices included in the framework resonated with the survey respondents: this translated into a validation of these specific activities and practices. In some other cases, novel Big Data activities and Data Governance practices emerged from the survey as potential additions to the governance standard. These additions to the framework require further validation, which is done through a group of experts on the subjects of Data Governance and Digital Ethics. Furthermore, the survey investigated how well commercial companies are currently dealing with Big Data ethics, and whether they have practices in place to limit the negative impact of Big Data handling on the data subjects. The validation session panel was confronted with the survey results and was asked to explain them.

6.1. Validation of framework

During the validation session, a panel of experts was walked through each section of the framework – corresponding to the data lifecycle phases – and was asked to comment on each addition to it, which was either a Big Data activity or a Data Governance practice. The experts asked, more specifically, to determine whether the presented activity or practice was valid, respectively from an ethical standpoint or from a Data Governance point of view; they were also asked to determine whether the relationship between the presented activity and practice was valid in the context of the framework. Furthermore, for each phase of the lifecycle the validation panel was asked to comment on the frequency of mentions of the corresponding governance practices, and to try to explain what those numbers meant in terms of assessing the status of commercial companies in addressing Big Data ethics. The following sections will go through the validation of the changes to the framework following the survey analysis, categorized by lifecycle phase (ordered by priority to keep the structure of the framework intact).

6.1.1. Data Collection

To begin with, the validation panel was asked to comment on the following Data Governance practice (highlighted in green) referring to the Data Collection phase:

Big Data Activity	Data Governance practice
Data collectors collecting data in illegal ways, for example by not asking users for consent in a specific, informed and unambiguous way, and by not knowing for what purpose the data itself will be used.	Applying GDPR requirements, involving data collectors making sure they have the right to collect personal user data, and that the users are aware of the scope and methods of data collection.

One Data Governance expert did not consider the formulation of the practice appropriate because it failed to answer the *How* question. Another expert stated that applying GDPR requirements is not something that can be actively done, and there is a whole compliance framework that explains how such requirements can be applied. It can be therefore considered more of a principle than a practice. They suggested that the practice could be reframed as “Implementing a control framework that applies GDPR requirements, involving data collectors making sure they have the right to collect personal user data, and that the users are aware of the scope and methods of data collection.”, and that this reformulated practice could be a good one for companies to undertake. It was also mentioned that the control framework applies to GDPR requirements, but it could also apply to other equivalents of that which exist in other countries.

For what concerns the corresponding Big Data activity, one digital ethics expert noted that abiding to laws is not equivalent to being ethical. From an ethical perspective, it is important that the data collectors do not collect data in an disproportionate way (which is what the GDPR tries to enshrine by making sure the data collected has a specific purpose, it is collected with the user consent etc.) and representative of all groups of people. It was clarified that this particular activity is focused on the legal matter of data collectors not asking for consent from the users, whereas the representativeness problem is addressed by a separate activity. As a consequence, the activity was not rephrased. The panel of experts confirmed then that the relationship between the presented activity and practice is valid.

The discussion moved to a second activity that needed validation:

Big Data Activity	Data Governance practice
Data gathering without limitation causes companies to not be aware of what data is being collected, with consequences on the end user.	Auditing the types of data collected and how it is secured.

An expert addressed this practice as a second line of defense. The group agreed that this practice is valid and it addresses the corresponding activity.

The third point that was validated is:

Big Data Activity	Data Governance practice
Sharing data within and outside a company in a non-transparent way, causing the 'opt-in' to use the data to be unknown by the company acquiring the data and causing the user to lose control of its personal data.	Setting up a data sharing agreement to govern the exchange of data between parties.

The content of the data sharing agreement is not described in this practice, so an expert proposed that the agreement should include a privacy impact assessment, which goes into different use cases and assesses the impact that data has based on what it is being used for. The practice is consequently rephrased to “Setting up a data sharing agreement to govern the exchange of data between parties, which includes a privacy impact assessment describing different use cases for the data being shared.”.

The last activity put through the lens of the panel for this lifecycle phase is:

Big Data Activity	Data Governance practice
Collecting personal identifiable information that violates people's privacy can affect customer trust - especially when uncommon types of data outside of customer expectations are collected - and the company's reputation, and puts the company at risk of legal noncompliance. (Kroll, 2018)	Anonymising data sources, encrypting important identifiers and relinquishing or agglomerating certain data fields to guarantee the users' anonymity.

One expert mentioned that this practice is something that will most likely be done later in the data lifecycle – specifically during the cleansing of data. However, it is appropriate to mention it in the Data Collection phase in order to reduce the risk of accidentally using the original dataset in the analysis. This practice was validated without needing to make changes to it.

6.1.2. Data Analysis

Moving only the Data Analysis phase, the first practice that the panel was asked to validate is:

Big Data Activity	Data Governance practice
Errors in predictions inferred from flawed Big Data models (e.g. biased AI or analytics models, or	Using a diverse analyst community (cultural, gender and ethnical diversity) to ease the removal of bias.

models having outdated data as an input) putting the safety of users in danger.	
---	--

The experts agreed that they were not sure it might be the best way to approach the corresponding activity. The focus should instead be on training employees, and on making sure they stay current with practices and procedures. The consensus was on removing this practice from the model, backed by the argument that having a diverse analyst community, while being something that in general should be done, will not necessarily be better for the outcomes of the model and cannot be considered a practice. Therefore this has to be replaced by a more concrete practice that specifically addresses the Big Data activity.

It was suggested that retraining models every once in a while would be a good solution for the activity: this would ensure that predictions are still correct over time. However, it is not always a feasible option. The practice "Introduce standardized operations to evaluate algorithms" was instead deemed to be exactly capable of addressing the activity and reduce errors in the predictions. Consequently, this practice was associated to the activity under consideration. Furthermore, an expert brought up the possibility of using specific tools to test and increase the transparency of algorithms, by checking what variables and what inputs are impacting the outcome. This was selected as an additional practice to address the activity.

As for the formulation of the activity under consideration, the example of models having outdated data as an input overlaps with the collection phase. The potential for use of outdated data should be in fact dealt with in previous phases of the lifecycle. It was suggested to only keep the example of the use of biased AI or analytical models in the formulation of the Big Data activity.

The following practice to be analyzed is:

Big Data Activity	Data Governance practice
Use of algorithms to target individuals in a personalized way.	Appointing a Data Protection Officer responsible for enforcing the measures needed to analyse data ethically.

Similarly to a similar, previously discussed practice, the practice of appointing a Data Protection Officer was deemed to be too generic for the framework. Especially after the GDPR it is easy for companies to fall back on a compliance mindset rather than an ethics mindset. While it could be that it is within the responsibility Data Protection Officer to make sure that data is analyzed ethically, this responsibility cannot be thrown exclusively on the shoulders of one person: it should instead be a responsibility shared among an entire analytics team.

A more actionable practice would be for the Data Protection Officer to establish an ethical agenda where it is made clear why a company wants to use a certain algorithm, why they need to target a specific person or a specific group, whether what the company does with data reflects the expectations of the customer and whether the act of analyzing data is good for the customer. By doing so, the company would be practically adopting an ethical approach to data rather than a compliance one.

The last part of the data analysis phase subject to the analysis of the validation panel is a combination of two practices to address the same activity:

Big Data Activity	Data Governance practice
Humans designing algorithms with their own perspective in mind introducing errors and biases that can harm users.	Training of professionals to be aware of the ethical and legal issues of data analysis.
	Coding reviews among data scientist to question potential ethical concerns.

For the presented practices, the panel agreed that they were actionable enough to not require any changes, and that they help address the corresponding Big Data activity.

6.1.3. Decision Making

For what concerns the Decision Making phase, no practice or activity required validation, since the results from the survey overlapped with the data previously collected from the literature and the expert interviews. However, by looking at this section of the framework, the experts noticed that the practice *Validate assumptions baked into the normalization methodology* is equivalent to the practice *Look for systematic biases in the way data is cleansed and validate data cleansing assumptions*: they suggested that the first practice should be merged under the second. The resulting practice is adequate to address both the activities of *Data scientists affecting data quality with the intention to get more interesting results* and *Data scientists choosing how to describe data and missing details in the world might introduce bias in the data*.

6.1.4. Data Cleansing

The part of the Data Cleansing phase that required validation is a combination of a novel Big Data activity and corresponding Data Governance practice proposed by a survey respondent:

Big Data Activity	Data Governance practice
Associating data sets together can reveal new, sensitive and confidential data that violates the user's privacy.	Aligning and validating the data cleansing activities with the data protection officer.

There was a debate regarding whether the Big Data activity would be part of the cleansing phase or the analysis phase. Given the definition of what activities are included in the data cleansing phase – which is the cleaning up of a database from empty or mistaken values, the activity under consideration should be moved to the Data Analysis phase within the framework. Assuming to have made this shift, the panel was asked whether the corresponding Data Governance practice is appropriate to address the activity. The discussion regarding this practice resembled what was previously discussed for the other similar practices within the framework: one expert stated that it is quite dangerous for a company to put too many responsibility on the Data Protection Officer, because it implies taking responsibilities away from individual data analysts – who should instead be having individual awareness of what is ethical to do with data and what is not. The panel agreed on the fact that validating the data analysis assumptions is a practice that would address the Big Data activity under consideration: this governance practice was therefore associated to the activity of associating data sets together in the framework.

6.1.5. Data Visualization

The part of the framework that needed validation for what concerns the Data Visualization phase is:

Big Data Activity	Data Governance practice
Exposing the visualized, sensitive customer data related to users to employees (especially to those who should not see it) and showing it in reports.	Anonymizing data before it is presented to the employees.
	Drafting guidelines that regulate which data can be shown to which employees in the company and which practitioners are mandated to work with customer-specific data.

The panel agreed with the formulation of both the Big Data Activity and the two corresponding practices. They however specified for the second practice that drafting guidelines is not by itself sufficient to address the activity, but the implementation of these guidelines is more crucial: only drafting such guidelines might increase the risk of companies treating this practice as a checklist and not putting enough effort into their actual implementation.

6.1.6. Data Archival and Data Deletion

For the Data Archival and Data Deletion phases there was one activity to validate:

Big Data Activity	Data Governance practice
Storing personal data past its due date with no clear business use.	Establishing a clear data retention policy that clarifies how and why data must be retained and how it will be used, to know how it can be profitably discarded and properly anonymized to minimize risk and make sure the retain process has no loose ends.

The panel of experts agreed with the Big Data activity to validate, but also added that in this section of the framework there is not enough attention given to the unethical deletion of data, which questions whether a company is deleting data without a proper call for deleting it. In fact, companies may need to keep some data for benchmarking purposes, to justify their decisions or to retrain their models. Sometimes they even need to keep some data to know which data they are not allowed to retain and collect.

Consequently, the unethical Big Data activity proposed is the preemptive deletion of data – the deletion of data before a company should actually delete it. The Data Governance practice that addresses it is similar to the practice proposed for the unethical retention of data: it is the formulation of a data deletion policy that describes what kind of data should not be deleted.

6.2. Validation of commercial companies' ethical data status

The validation session also served to explain the results of the survey concerning the evaluation of how ethical commercial companies are in handling personal data. This evaluation is supported by survey data regarding the frequency of usage of governance practices to address Big Data-related ethical problems, as well as the priority and importance attributed to specific Big Data lifecycle phases.

6.2.1. Frequency of mentions of Data Governance practices

The panel was walked through the phases of the data lifecycle and showed the frequency of mentions of the governance practices concerning each of these phases. They were then asked to explain what it can say about the data processes of the surveyed companies.

Starting from the Data Collection phase, one explanation for why the first four practices got the most mentions is that they are formulated in a way that they are more easily defined and recognizable as traditional data governance practices. *Reduce error of data* is more of a practical endeavor rather than a policy, procedure or strategy. *Define a data strategy* is a very wide practice that incorporates most of the other practices and is quite high level. Two experts agreed that they were surprised to see the practice *Design a review board with the power to approve or deny the collection of new data* low in the list (with only three respondents mentioning they have it in place within their organizations): in practice this is not so surprising given that there are not many review boards around within companies, however they expect this practice to appear higher in the list in the future.

They also noted that the overview of the used practices helps explain the maturity of the surveyed companies in addressing Big Data ethics. In fact, within an organization wanting to address ethics the first step would be to set up the *accountability for the data collection*: however, there is a further step required after determining who is accountable for this phase, which is to act upon it by, for example, actively anonymizing data sources and encrypting important identifiers. Understanding that these activities need to be undertaken comes with time, so in a second moment after determining who is accountable for the data collection.

The fact that execution-type of practices such as *Reduce error of data* appear so low in the list of mentions can be a sign of maturity. From the results, it looks like commercial companies are in a less mature stage where

they discuss ethics in a more conceptual way, in terms of defining accountabilities and a data strategy. The results obtained therefore suggest a lack of ethical data handling maturity of the surveyed commercial companies.

The list of practices mentioned by the survey respondents related to the Data Analysis phase shows some similarities to what was previously seen for the Data Collection phase. In fact, the practices mentioned most times for this phase are once again vague and high level: assigning a data owner, determining who is accountable for the data usage, establishing policies and procedures or defining a data strategy are practices show that companies are not thinking of ethical questions at a deep level. While these high level practices do constitute a good foundation for engaging in ethical data handling, they are in a sense equivalent to checking boxes: these practices are not effectively contributing to executing Big Data activities more ethically. From the survey results we can understand that companies are immature in the sense that they are still in an exploratory phase of how they can use new technologies and algorithms and how they can get the most value out of them, and are not yet in the stage of asking themselves what the ethical consequences for using Big Data are. This would explain why the surveyed companies wouldn't want to engage in the practice *Introduce standardized operations to evaluate algorithms* – which is low in the list of mentioned practices: introducing standardized operations would in fact limit the exploration of the new Big Data technologies. It is noted by an expert that from an ethical perspective this would be a more effective practice to undertake than establishing Accountability of data usage - which is a very vague practice especially when discussing very concrete activities such as training models and doing analysis on data: if a company wanted to validate their ethicalness in handling data, there should be a standardized way to evaluate the algorithms every several months. Differently to what seen in the previous phase, the practice of *Designating a review board* is higher in the list. This might be justified by thinking that the Data Analysis phase is seen by companies as a more active phase, and it is easier to have more control over how data is used in comparison to where and how you are getting the data.

For what concerns the Decision Making phase, the results showed to the validation panel brought up the observation that decision making is still done in an exploratory way within organizations, where the employees would have a graph and would make decisions based on the data represented on the graph instead of having an algorithm making decisions for them. Companies are still exploring new technologies instead of engaging in practices that monitor how certain results are generated: this is where companies might fall foul of ethical considerations. The practice *Monitor the performance of a system after launch by means of black box* is the most actionable and effective practice among the ones in the list, because it is about running options at the same time to see which one has the most fallout or which will result in the most business benefits. The fact that only one survey respondent stated they have this practice in place is a sign that companies are not there yet in terms of being mature in how they make data-based decisions.

The mentions of practices related to the Data Cleansing phase are a confirmation of a lack of maturity of the surveyed companies: having more actionable practices on top of the list would be desirable from an ethical perspective, but what the results show is that the surveyed companies do not engage in such practices such as the *Validation of data cleansing activities* and the *Validation of data cleansing assumptions*. They are rather occupied with defining a data strategy or establishing policies and procedures.

The results concerning the Data Visualization phase were surprising compared to the other phases: the practices that got mentioned more frequently were in fact more actionable than the most selected practices for the previous phases – which were on the contrary more vague and high level. This difference could be explained by noting that the Big Data activity of exposing sensitive data to employees has a very tangible damage and therefore you would need to have very tangible controls against it. The unethical activities mentioned in the other phases are instead perhaps more abstract and amorphous, so the responses for them are consequently also more abstract and amorphous. It was added that the problem of employees being exposed to sensitive data – sometimes even data concerning themselves or their colleagues, has existed for a long time and it is a problem that lives within the type of people that deal with visualizing data. Because these people are used to facing this problem, the practices to address it take a more concrete and actionable form.

Lastly, looking at the results for the Data Archival and Data Retention phase, it was noted that the practice of *Establishing a clear data retention policy* is very complete and therefore it makes sense for the surveyed companies to have it in place and have selected it most frequently.

6.2.1. Perceived importance and priority level of the Big Data lifecycle phases

The second way that the survey addressed the ethical status of commercial companies in how they handle personal data was by asking them to rank the Big Data lifecycle phases in terms of the priority each of them has within the organization, and how important they are for the survey respondents. The experts were asked to comment on these results, but first they were engaged in an exercise where they had to come up with a list of criteria they would use to make a ranking of the data lifecycle phases; they then had to apply such criteria to determine which phases of the lifecycle are more important for commercial organizations to address due to their ethical riskiness level.

One expert mentioned that, when determining which criteria he would use to make the evaluation, he would consider which phases of the lifecycle involve acting upon data or information generated. Consequently, he made the assessment that the decision making phase is the lifecycle phase that brings the most risk to a company because, ethically speaking, the issues start if decisions are based on data, so when action is taken on the basis of the data that has been collected, analyzed and visualized. They argued that there isn't an ethical issue if nothing comes out of the information generated by an algorithm.

“Having information on something you shouldn't have known is not as bad as taking an action based on the information you shouldn't have had in the first place: there is a stage in between where a human would ethically judge that it is not ethical to make a decision based on the information they have.”

For example, it wouldn't be wrong to know that certain people behave in a certain way: however issues arise if, out of that information, a company decided to push advertisements to those people knowing they might be susceptible to certain things.

There was some debate regarding this, however: another expert stated that deciding not to collect risky data would put a company in an easier position in the decision making phase, and he therefore thought that the data collection phase is the one that brings the most risk for companies. One panelist argued instead that he would use as a criteria the impact that a certain phase can have on end users. He believes that ethical debates are about trust, and trust bows down to a discrepancy between expectations and reality. The ethical risk in using Big Data is, to him, losing the trust of customers, therefore one should consider which activities carry this risk to make an assessment of the lifecycle. When the customers' expectations about what a company will do with their data does not match the reality of facts, when customer data is used against them, when the data behaviors clash with the values of the customers, companies face the highest ethical risks.

The purpose of this exercise was not so much to get a right answer, but rather to show that the criteria that one can use to make an assessment of importance of the lifecycle phases can vary from person to person. The criteria reflects different perspectives that one can choose to be adopt, and could be the impact on users, the possibility of losing customer trust when their expectations don't match the data behavior of companies, or which phase involves acting upon information.

The validation panel was at this point showed the results concerning the level of priority that each lifecycle phase has within the surveyed organizations. Supposedly, this data should be saying something about where these companies stand in addressing Big Data ethics. One expert commented that the list reflects where the organizations currently see the most risk, and the reason why the validation phase occupies the higher position might be that it is the phase where ethical violations are found out: this phase is, in fact, the culmination of a series of activities individually bringing unethicalness which then appears more prominently when the outcome of a model is visualized. The lifecycle phase priority from the survey respondents' side show the adoption of an inside-out perspective. When adopting this perspective, the organization feels they run the most risk with the visualization of data because that is the face where they see more directly the impact of the data analysis results: they see the risk that their employees see information that they are not

supposed to see, as well as the risk to have data they are not supposed to have because they did not archive it or delete it correctly. Consistently with this, the decision making phase was ranked last by the survey respondents, likely because this phase requires an outward perspective on the customer and how the decisions made might impact them.

It was also explained that the reason why the responses of the survey participants differ from the expert opinions is that the two groups adopted a different perspective when making the evaluation of the lifecycle. While the survey respondents adopted an inside-out perspective, it is likely that the experts adopted a customer-focused approach to the lifecycle, which explains why the data collection phase is at the top of their priority. This reasoning is also in line with the exercise that the validation panel did previously: in fact, the panel did agree on the fact that the survey respondents used different criteria than the experts to rank the lifecycle phases. A similar reasoning applies when comparing the perceived importance of the lifecycle phases by the survey respondents with the expert opinions. Also in this case, a difference in perspectives adopted when judging the lifecycle phases can be noticed, which is evident once again from the fact that the visualization phase is high in the list and the decision making phase is rated last – consistently with what discussed for the results concerning the lifecycle phases priority list.

The results, however, don't necessarily imply a lack of maturity of the surveyed organizations, but it does clarify that the approach that the surveyed commercial companies use is inwards looking: these companies go through the phases of the lifecycle driven by a will to maximize the business value of using Big Data and its related technologies, and get to the decision making point without considering what impact it will have on the customer, what value it will bring them and whether the business decisions match the customer expectations. The results overall imply a heavy prioritization and focus on internal capabilities rather than values, where the technical people handling data are let loose and the organization lacks a clear, holistic approach to data ethics.

7. Discussion

Despite the risk that the collection and use of Big Data poses to user privacy, existing research shows that no practical indication of how commercial companies can mitigate such risks. Data Governance has been mentioned as a way to establish an ethical culture within an organization (DAMA International, 2017), however it is not clear which controls should be introduced to address Big Data-related risks and thus ensure that the processing of data does not violate any ethical principle, human right and data regulation.

Therefore, the aim of this study is to understand why it is important for organizations to address the topic of ethics in their handling of personal user data, and whether the concept of Data Governance can be translated into concrete practices that can support such organizations in their ethical journey. The results of the study can be used to argue whether the concept of Big Data ethics can be operationalized and used to determine governance practices aimed at addressing the risks the unethical handling of data. Secondly, the results allow the researcher to argue how well-prepared commercial organizations are at tackling such risks – and therefore how ethical their data processes are.

A combination of qualitative and quantitative methods was used to answer the main research question of this study. The literature review was used to define and operationalize the concept of Big Data ethics, and to identify which activities involving the use of Big Data are known to carry ethical risks. A round of expert interviews served to explore such activities further and understand the motivation behind classifying them as unethical; the expert opinions also contributed to the development of a Data Governance framework that can aid organizations in carrying out an ethical transformation of their data processes by tackling risky, unethical Big Data activities with corresponding governance practices. Lastly, a group of data practitioners was surveyed to investigate whether the governance practices are already in use within commercial companies: this helped create a picture of the level of readiness of such organizations in addressing the ethical risks that come with handling personal data.

7.1. Usability of Data Governance framework for Big Data ethics

The survey and the validation session were able to qualitatively confirm that a relationship exists between Big Data activities and Data Governance practices. In fact, throughout the survey the respondents expressed that governance practices exist within their companies that address risky, unethical activities involving Big Data. During the validation session these relationships were validated, confirming that the adoption of governance practices within an enterprise could limit the impact that unethical activities may have on a company's reputation and on their customers. The Data Governance framework for Big Data ethics makes these relationships explicit, by showing which governance practices are suitable to address which Big Data activities.

In terms of applicability, the framework is an instrument that could be used by companies looking for ways to deal with Big Data activities in an ethical way, providing them with a set of hands-on, actionable practices – some more generic, some more detailed, ordered by lifecycle phases that prioritize the impact on the customer.

Regarding usability, what resulted from the validation session is that the framework could be used to determine what kind of perspective a certain organization adopts when managing their data process. Despite the fact that the validation panel argued that the framework is not able to make a judgement on the level of maturity of an organization when it comes to addressing Big Data ethics, one could argue that the adoption of an inward or outward perspective can be an indicator of the ethical maturity of a company. Specifically, the adoption of an inward perspective which does not focus on the needs of the customers and on respecting their privacy, is a symptom of ethical immaturity: in fact in order to be ethical an organization should organize their data processes in a way that the privacy and human rights of the users are safeguarded. By not putting their focus on the customer, organizations are showing a lack of ethical maturity. On the other hand, prioritizing the safeguarding of human rights and individual privacy would cause a company to not engage in unethical activities that would put those things in danger, thus implying a higher ethical maturity in the way they handle

user data. The framework can therefore be used as a checklist to determine how mature a company is in terms of data ethics. By using the framework in such a way, an organization should be able to determine how well prepared they are in dealing with the ethical consequences of personal data handling – based on how many activities and practices they can check off the list. This shows the framework’s ability to help organizations assess how they deal with Big Data ethics. The framework is also normative in that it could instruct companies on what to do if they were to engage in certain unethical activities involving the use of Big Data. There is a scenario in which a company, by looking at the framework, realizes that they engage in many unethical activities: this could be a sign of an approach to Big Data ethics which isn’t enough ethics-centered. In this case the enterprise might want to organize its Data Governance by adopting the practices recommended by the framework. A different scenario is one in which an organization does not recognize the Big Data activities in the framework as activities they engage in, which could be a sign that they are well prepared for dealing with Big Data ethics – or at least better prepared than those companies who do experience such activities. It should be noted that one parameter to consider when using the framework is in which context the organization uses Big Data: the risky activities they might engage in might depend on the particular use they make of Big Data (e.g. a company might be centered around the data analysis more than its collection). Therefore, framework users should not be comparing the number of unethical activities they engage in against each other, but also consider the context they work in.

Assuming the head of the data office of an organization has pointed out that they want their company to become more ethical, data officers and data protection officers under their umbrella may take lead roles and use the framework to check whether the organization engages in unethical Big Data activities – and if they are, check if they are including a data governance practice to address them. The entire framework is therefore mainly addressed to those who are part of a data office and have an oversight over the data process of the whole company. People within the data office should take the practices of the framework that suit their company best and bring them to their employees. Employees who have a more narrow focus on the data lifecycle, such as data scientists, might find it difficult to make use of the framework themselves without an incentive from the strategy level of the company: in fact, following the guidelines of the framework might restrict their day-to-day activities. However, if their managers wanted to implement ethical practices within the organisation, they could translate such practices into ethical day-to-day activities for the data scientists to follow. The approach of utilising the framework is therefore a top-bottom approach, especially considering the current scenario of ethical immaturity that companies are in. The ethical awareness and consequent practices should come from the head of the data office and be integrated into existing processes spanning the whole data lifecycle. Ideally, the approach should turn over time into a bottom-up approach, where ethics are integrated in a company in activities such as the hiring process and the talent reviews.

6.2. Status of commercial companies in addressing Big Data ethics

By means of a survey, a group of data practitioners working in commercial companies was questioned about the ethicalness of the data processes of their respective organizations. An ethical assessment of the surveyed sample was made by combining data of the frequency of usage of Data Governance practices (aimed at addressing unethical Big Data activities) and the ranking of the data lifecycle phases based on the priority level and perceived importance associated to each phase.

What emerged is that the population of commercial companies surveyed tends to adopt an inside-out thinking that makes them think of Big Data ethics from an inside perspective, prioritizing their own points of view and their own responsibilities. The effect of this inward thinking when looking at the Big Data lifecycle is a shift away from acting in the customer’s best interest. By contrast, adopting an outward perspective would mean looking at the lifecycle from a customer point of view, and acting on the company’s data processes based on an evaluation of the consequences that certain activities would have on the customer’s rights such as data privacy. The group of experts interviewed prior to the survey have shown to adopt this outward perspective when judging the data lifecycle, possibly due to their higher consideration they have towards the impact of Big Data on customers rather than on a specific company.

The results of this study might be a reflection of the industry lacking a holistic knowledge of the whole data process. In fact, when comparing what the surveyed companies do and what they think is correct to do in the future in terms of prioritizing the lifecycle phases, the outcomes are very similar and are both inward-focused. Therefore, one could argue that these companies' vision of Big Data ethics is skewed and not heading in the right direction. However, companies are not the only party holding responsibility in this debate. These results could also be the reflection of the lack of an industry standard that acts as a reference for all data practitioners and enforces an ethical perspective on the Big Data lifecycle of activities. The survey showed that there is clearly awareness of the problem of Big Data ethics among data practitioners, however these don't have the same outcome in mind as the experts do.

8. Conclusion

In this section the results of the study will be synthesized and the conclusions of the study will be derived from them. The research questions formulated at the beginning of the research process will be used as a guideline to present the conclusions; furthermore, going through each research question will help determine whether the study's goals have been reached.

[1] What are Big Data ethics?

A review of existing literature highlighted the lack of a comprehensive definition of Big Data ethics – which has so far been referred to in literature as data ethics. Previous definitions did not take into consideration the recent phenomena of digital information growing in volume, velocity and variety, as well as the rise of technologies that allowed companies to turn this information into useful insights by means of Artificial Intelligence, Machine Learning and Data Mining. The increased risk of business ethics violations connected to such new opportunities (Herschel & Miori, 2017) raised the need of introducing a novel field which studies the impact that collecting and processing Big Data, as well as making decisions on the basis of the derived insights, may have on the end user. Consequently, this study introduced the concept of Big Data ethics, which has been defined as follows:

“Big Data ethics is a branch of business ethics that studies ethical problems that arise in the business environment when using Big Data and algorithms for data analysis. Its goal is to develop moral rules, standards, or practices that support moral decision-making based on Big Data analytics”.

[2] What is the role of existing data laws and regulations in addressing Big Data ethics?

In 2017 the GDPR regulation on data privacy was introduced to give users more control over their data, bringing increased awareness to the topic of data ethics. Data laws and regulations should be seen as a starting point to develop an ethical culture within an enterprise, however literature and expert interviews brought light to the fact that organizations are merely trying to be compliant to the GDPR principles. What could explain such behavior is the fact that organizations do not want to disrupt the way they usually conduct business more than necessary, and they do not see the added value of integrating ethics into their day-to-day data processes. However, handling customer data unethically can have dangerous consequences, and in order for organizations to truly safeguard themselves from them they should work to develop a constantly evolving strategy for data protection and privacy, one that would enable them to respond to evolving technologies (ibe, 2018).

[3] How do fundamental ethical principles relate to the Big Data context?

Ethical principles are rules that help guiding one's behavior and can support them in solving ethical dilemmas in everyday life (Weinstein, 2017). The most established ethical principles that have been defined in literature are the principles of Respect for Autonomy, Maleficence, Beneficence, Justice and Privacy and Data Protection. These principles, however, would be difficult for organizations to apply as they are formulated; thus, an interpretation of how they apply in the specific context of Big Data is required in order for organizations to be able to use them to mitigate the risks of Big Data handling. Consequently, the researcher took the task of establishing a relation between such ethical principles and Big Data by operationalizing the ethical principles into ethical values – which can help identify which actions are risky and may impact the end user negatively. On the basis of such values, this study has identified a list of examples of how the collection of personal user data, its processing and the use of the derived insights to make business decisions can infringe data regulations and human rights. This list of unethical activities is the result of the application of ethical principles to the context of Big Data and constitutes the basis of a governance standard that aims to tackle such unethical activities.

[4a] What Data Governance practices are currently being used by commercial organizations to address Big Data ethics and data regulations?

What resulted from conducting a survey among data practitioners working within commercial organizations is that, for each phase of the lifecycle, the Data Governance practices they have in place to mitigate the ethical risks that come with Big Data are quite strategic and high-level, and abstract in the way they are formulated. Such practices include, for example, the definition of a data strategy and the establishment of policies and practices. Based on these results, we can deduce there is a lack of ethical maturity within commercial organizations in the way they collect, process and make decisions based on the analysis of Big Data, and a bigger will to reap the benefits of Big Data analytics and emerging technologies than to do business in respect of the customer's rights.

From the expert's point of view, a data process that puts the focus on the customer is required to ensure ethical behavior; however, what emerged from the survey analysis is that commercial companies are internally focused and prioritizing compliance to data regulations, rather than how their daily operations affect the customer. This shows a difference between practice and theory and what each party prioritizes. From these results we can conclude that commercial organizations have a skewed focus which looks at the Big Data lifecycle while only having the interests of the company in mind, instead of looking at the potential impact of Big Data-based decisions on the customer.

[4b] What Data Governance practices should be used by commercial organizations to address Big Data ethics and data regulations?

When it comes to using Data Governance to make the Big Data lifecycle more ethical overall, the researcher agrees with the experts of the validation panel in that a holistic approach should be used – one that looks at the lifecycle as a whole and integrates ethical considerations in each of its phases. Adopting governance practice for single phases of the lifecycle is likely to not yield satisfactory results and definitely safeguard an organization from ethical violations, even when such phases being addressed are the most risky from a customer perspective: for example, adopting practices that would ensure data is being collected in respect of the customer's rights would not necessarily prevent an organization from making unethical decisions with such data at later stages of the lifecycle. Furthermore, the specific practices that should be adopted by commercial organizations vary depending on their current level of ethical maturity. For an organization that has only just determined they want to transform their data processes into one that considers the ethical impact on the customer, this study suggests, as a starting point, to use Data Governance practices that are more generic in nature: this could be the definition of a data strategy or the drafting of a control framework that clarifies how certain actions are to be executed within the Big Data lifecycle. More mature companies that are already on the path of making their data processes more ethical should instead aim to translate a data strategy into practical, day-to-day actions, to be implemented from the data office down to each employee that handles data at some stage of the data lifecycle.

How can Data Governance support commercial organizations in addressing Big Data ethics?

Existing literature does not clarify how a Data Governance program would enable organizations to reap the benefits of Big Data, while tackling the ethical risks that Big Data brings. In this study, a list of unethical Big Data activities and Data Governance practices aimed to address such activities were organized into a Data Governance standard for Big Data ethics. The developed framework (Annex J) proves that Data Governance can be instrumental in driving the journey of commercial organizations towards a more ethical handling of personal user data. The standard consists of guidelines that commercial organizations can put into practice to limit the negative impact of unethical actions they might engage in at a certain stage of the Big Data lifecycle. The structure of the framework reflects the opinions of the interviewed experts, who put emphasis on the riskier phases of the lifecycle. However, the researcher believes that the proposed governance practices are to be integrated in an organization's data process holistically throughout the data lifecycle. The framework also prioritizes the lifecycle phases that have a higher impact on the customer: this is coherent with the results of the survey analysis, which showed that a customer-centric approach is required by commercial organizations in order to achieve ethical maturity. Such approach should consider the customer's interests when managing

the data activities throughout the lifecycle.

It should be noted that not all recommended practices may be suitable for every organization, as the scope of the use of Big Data within their operations may vary. With that considered, the framework's objective is primarily to be a starting point for any ethical discussion concerning the use of Big Data within commercial organizations. The presented practices can support commercial companies in their efforts to organize and transform their data processes in such a way that the rights of their customers are not violated, and that they comply to data regulations such as the GDPR.

8.1. Research implications

The results agree with previous research which identifies Data Governance as an instrument to integrate ethics into the data processes of an organization. Existing literature however focused on introducing the role of Data Governance into the data ethics debate, rather than explaining how organizations should put it into practice. The results of this study demonstrated that by investigating where ethical violations may occur throughout the Big Data lifecycle, it is possible to identify very specific governance practices that can mitigate the risks of such activities.

The data also contributes a clearer understanding of the current status of commercial organizations in addressing Big Data ethics: these were shown to be ethically immature. Compliance to new data regulations such as the GDPR have been a strong motive for organizations to start a Data Governance program; the latter, however, does not seem to have been translated into ethical actions at an operational level within the surveyed commercial companies. This highlights the need of an industry standard that can act as a reference for organizations that wish to handle personal data more ethically – not only for compliance reasons. In fact, this study has also shown that the ethical handling of user data could be a selling point for organizations, and that the integration of ethics into Big Data could be transformed into a competitive advantage and ultimately encourage new business models to arise. These results should therefore be taken into consideration by those organizations struggling to come up with a business case for data ethics.

8.2. Limitations

For what concerns the construction of the Data Governance standard for Big Data ethics, this has been built solely on the basis of qualitative data; this means that the quality of the framework itself is dependent on the level of expertise of the study participants that supported its construction – namely the experts of the first round of interviews and the survey respondents. To mitigate the risks of bias of the study participants, for the interviews and the survey a diverse sample of experts and practitioners was selected: this diversity is reflected in their job roles, which span across the fields of law, data privacy and data science, as well as the years of experience in their respective fields. Furthermore, the framework was built on top of existing literature and went through several rounds of validation. It is beyond the scope of this study to test the framework in a practical setting, as this may require modifications to suit the specific organizational context it is applied in.

As for the investigation of the ethical status of commercial companies in addressing Big Data ethics, it should be noted that the generalizability of these results is impacted by the relatively small sample size of 31 organizations who participated in this study's survey, as well as by the bias potentially introduced during the validation panel that supported the explanation of the survey outcomes. Moreover, this study is not able to provide an overview of how ethically mature a specific industry is, being that the distribution of industries in the survey sample is not even.

8.3. Future research

Future research should look into the practical applicability of the Data Governance standard for Big Data ethics within the context of commercial organizations, as this was not tested in this study. Following studies should also consider investigating whether it is suitable to use for non-commercial companies, such as the private sector or the hospitality industry, and generalize the framework if necessary.

The framework could also be expanded to include an element that defines critical roles and responsibilities associated with the process of integrating ethics into the Big Data lifecycle of an organization. Potentially, such roles could be independent from the data team so that the framework could be used as a tool for the internal auditing of a company's data process.

As for a practical follow-up to this study, the proposed framework could be used as a guideline for the development of a generally recognisable industry standard for Big Data ethics.

Furthermore, a result of the study was that the perspectives on the Big Data lifecycle are different between the population survey respondents and the experts. Future research might go more in depth into this and understand the extent to which the priorities might differ among data-related roles within a single organization. For example, a data scientist might have a different view on the data process than managers, which means they will put their focus on different lifecycle phases. By proving that different perspectives exist and are adopted within a single organizations, the discussion could move towards understanding how this gap can be closed so that the organization can embark on their ethical journeys in a more unite way.

Lastly, in order to increase the usability of the framework further, future research could use the developed Data Governance standard as a starting point to create a maturity model, to be used by commercial companies to assess where they stand in terms of Big Data ethics and to provide them with concrete points they should follow in order to move up to a higher level of maturity.

Bibliography

- ACCA. (2014, November 25). *Why is ethics important to business?* Retrieved from ACCA Blogs: <https://blogs.accaglobal.com/2014/11/25/why-is-ethics-important-to-business/>
- Adom, D., Joe, A. A., & Hussein, E. K. (2018, January). Theoretical and Conceptual Framework: Mandatory Ingredients of a Quality Research. *International Journal of Scientific Research*, 7(1).
- Alshboul, Y., Wang, Y., & Nepali, R. K. (2015). Big Data LifeCycle: Threats and Security Model. *Twenty-first Americas Conference on Information Systems*. Puerto Rico.
- Alshenqeeti, H. (2014). Interviewing as a Data Collection Method: A Critical Review. *English Linguistics Research*, 3(1).
- Amakobe, D. F. (2016, October). *Business Ethics*. Retrieved from Research Gate: <https://www.researchgate.net/publication/308926602>
- Asadi Someh, I., Breidbach, C., & Davern, M. (2016). Ethical Implications of Big Data Analytics. *Twenty-Fourth European Conference on Information Systems (ECIS)*. Istanbul, Turkey.
- Asadi Someh, I., Breidbach, C., & Davern, M. (2016). Ethical Implications of Big Data Analytics. *Twenty-Fourth European Conference on Information Systems (ECIS)*. Istanbul, Turkey.
- Baeza-Yates, R. (2013, January). *Big Data or Right Data?* Retrieved from Semantic Scholar: <https://pdfs.semanticscholar.org/4b94/5ee77794895454447ed5eb80c65f6c974a74.pdf>
- Beauchamp, T. L., & Childress, J. F. (2001). *Principles of Biomedical Ethics*. New York: Oxford University Press.
- Bhadani, A. K., & Jothimani, D. (2016). Big Data: Challenges, Opportunities and Realities. In M. K. Singh, & K. G. Dileep, *Effective Big Data Management and Opportunities for Implementation* (pp. 1-24). Pennsylvania, USA: IGI Global.
- Boyd, D., & Crawford, K. (2012, June). Critical Questions for Big Data. *Information, Communication & Society*, 15(5), 662-679.
- Burnard, P. (1999). Searching for meaning: a method of analysing interview transcripts with a personal computer. *Nurse Education Today*, 14, 111-117.
- Business Ethics Briefing. (2018, January). *Business Ethics and Artificial Intelligence*. Retrieved from ibe: https://www.ibe.org.uk/userassets/briefings/ibe_briefing_58_business_ethics_and_artificial_intelligence.pdf
- Cary, C., Wen, J. H., & Mahatanankoon, P. (2003). Data mining: Consumer privacy, ethical policy, and systems development practices. *Human Systems Management* 22, 157-168.
- Castillo-Montoya, M. (2016). Preparing for Interview Research: The Interview Protocol Refinement Framework. *The Qualitative Report*, 21(5), 811-831.
- Chessell, M. (2014). *Ethics for big data and analytics*. Retrieved from IBM: https://www.ibmbigdatahub.com/sites/default/files/whitepapers_reports_file/TCG%20Study%20Report%20-%20Ethics%20for%20BD%26A.pdf
- Clark, C. T. (2019, January 23). *The Ethics Of Data Governance - 'Data Comes With Benefits And Liabilities'*. Retrieved from Forbes: <https://www.forbes.com/sites/charlestowersclark/2019/01/23/the-ethics-of-data-governance-data-comes-with-benefits-and-liabilities/#3b157a68215a>
- Clarke, R. (2016, July 24). *Introduction to Dataveillance and Information Privacy, and Definitions of Terms*. Retrieved from Rogerclarke.com: <http://www.rogerclarke.com/DV/Intro.html#Priv>
- Crawford, M. M., & Wright, G. (2016). Delphi Method. In *Wiley StatsRef: Statistics Reference Online* (pp. 1-6).
- DAMA International. (2017). *DAMA-DMBOK Data Management Body of Knowledge*. Basking Ridge, New Jersey: Technics Publications.

- Davis, K., & Patterson, D. (2012). *Ethics of Big Data*. O'Reilly.
- De George, R. T. (1987). The Status of Business Ethics: Past and Future. *Journal of Business Ethics*, 6, 201-211.
- De Mauro, A., Greco, M., & Grimaldi, M. (2015). What is Big Data? A Consensual Definition and a Review of Key Research Topics. *International Conference on Integrated Information (IC-ININFO. 1644)*, pp. 97-104. AIP Publishing.
- Denscombe, M. (2010). *The Good Research Guide for Small-Scale Social Research Projects*. Butterworth-Heinemann.
- Duan, Y., Edwards, J. S., & Dwivedi, Y. K. (2019). Artificial intelligence for decision making in the era of Big Data – evolution, challenges and research agenda. *International Journal of Information Management*, 48, 63-71.
- El Arass, M., & Souissi, N. (2018). Data Lifecycle: From Big Data to Smart Data. *5TH EDITION INTERNATIONAL IEEE CONGRESS on INFORMATION SCIENCE and TECHNOLOGY (CiSt'18)*. Marrakech, Morocco.
- El Arass, M., & Souissi, N. (n.d.). Data Lifecycle: From Big Data to Smart Data. *5TH EDITION INTERNATIONAL IEEE CONGRESS on INFORMATION SCIENCE and TECHNOLOGY (CiSt'18)*. Marrakech, Morocco.
- European Commission. (2007, June 14). *Ageing well in the Information Society: Action Plan on Information and Communication Technologies and Ageing*. Retrieved from EUR-Lex: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=LEGISSUM%3AI24292>
- Farlex. (n.d.). *Business Practice*. Retrieved from The Free Dictionary by Farlex: <https://financial-dictionary.thefreedictionary.com/Business+Practice>
- Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996). From Data Mining to Knowledge Discovery in Databases. *American Association for Artificial Intelligence*. Retrieved from <https://www.kdnuggets.com/gpspubs/aimag-kdd-overview-1996-Fayyad.pdf>
- Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996, November). The KDD Process for Extracting Useful Knowledge from Volumes of Data.
- Fischer, J. (2004). Social Responsibility and Ethics: Clarifying the Concepts. *Journal of Business Ethics*, 52, 391-400.
- Floridi, L., & Taddeo, M. (2016). *What is data ethics?* Retrieved from Royal Society Publishing: <https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2016.0360>
- Foley, B. (2018, June 6). *Performing Qualitative Research with Surveys*. Retrieved from surveygizmo: <https://www.surveygizmo.com/resources/blog/performing-qualitative-research-with-surveys/>
- Gartner. (2018, November 29). *Data Ethics Enables Business Value*. Retrieved from Gartner.com: <https://www.gartner.com/document/3894128?ref=solrAll&refval=229927591&qid=87e8e5e0d4c6800b94ccf2b1>
- George, G., Haas, M. R., & Pentland, A. (2014). Big Data and Management. *Academy of Management Journal*, 57(2), 321-326.
- Gert, B. (1999). Common morality and computing. *Ethics and Information Technology*, 57-64.
- Granger, M.-P., & Irion, K. (2018). The right to protection of personal data: the new posterchild of European Union citizenship? In d. W.-P. de Vries S., *Civil Rights and EU Citizenship*. Cheltenham: Edward Elgar Pub.
- Grant, C., & Osanloo, A. (2014). Understanding, Selecting, and Integrating a Theoretical Framework in Dissertation Research: Creating the Blueprint of your "House". *Administrative Issues Journal: Connecting Education, Practice, and Research*, 12-26.

- Haasdijk, E. (2019, April 4). *Open the 'black box' of Artificial Intelligence with Glassbox (2/5)*. Retrieved from Deloitte: <https://www.deloitteforward.nl/artificial-intelligence/open-de-black-box-van-artificial-intelligence-met-glassbox/>
- Han, J., Kamber, M., & Pei, J. (2001). *Data Mining: Concepts and Techniques*. Morgan Kaufmann.
- Han, J., Kamber, M., & Pei, J. (2012). *Data Mining Concepts and Techniques*. Waltham: Morgan Kaufmann.
- Hänold, S., Forgó, N., Monreale, A., Ruggieri, S., van den Hoven, J., Mahieu, R., & van Putten, D. (2016, August 31). *Value-Sensitive Design & Privacy-by-Design technologies for big data analytics*. Retrieved from Cordis EU research results: <https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5acc4dff&appId=PPGMS>
- Harrie, J. (2010). *The Logic of Qualitative Survey Research and its Position in the Field of Social Research Methods*. Retrieved from Forum: Qualitative Social Research: <http://nbn-resolving.de/urn:nbn:de:0114-fqs1002110>
- Herschel, R., & Miori, V. M. (2017). Ethics & Big Data. *Technology in Society*(49), 31-36.
- Holsapple, C. W., & Joshi, K. D. (2002). Knowledge manipulation activities: results of a Delphi study. *Information & Management*, 39, 477-490.
- Horton, M. (2019, February 25). *Why is business ethics important?* . Retrieved from Investopedia: <https://www.investopedia.com/ask/answers/040815/why-are-business-ethics-important.asp>
- Hsu, C.-C., & Sandford, B. A. (2007, January). The Delphi Technique: Making Sense Of Consensus. *Practical assessment, research & evaluation*, 12(10), 1-8.
- ibe. (2018, May). *Beyond Law: Ethical Culture and GDPR*. Retrieved from Business Ethics Briefing: https://www.ibe.org.uk/userassets/briefings/ibe_briefing_62_beyond_law_ethical_culture_and_gdpr.pdf
- JEPS Bulletin. (2018, March 1). *Writing a Systematic Literature review*. Retrieved from JEPS Bulletin: <https://blog.efpsa.org/2018/01/03/writing-a-systematic-literature-review/>
- Koops, B.-J., Newell, B., Timan, T., Škorvanek, I., Chokrevski, T., & Galič, M. (2017). Typology of Privacy. *Legal Scholarship Repository*.
- Kramer, L. (2019, June 13). *Similarities Between Personal Ethics & Business Ethics*. Retrieved from bizfluent: <https://bizfluent.com/info-7742837-similarities-personal-ethics-business-ethics.html>
- Kroll, J. A. (2019, January 21). Data Science Data Governance. *AI Ethics*.
- Law v. Canada (Minister of Employment and Immigration). (1999). SCR.
- Lewis, P. V. (1985, October). Defining 'Business Ethics': Like Nailing Jello to a Wall. *Journal of Business Ethics*, 4(5), 377-383.
- Liu, J., Li, J., Li, W., & Wu, J. (2015). Rethinking big data: A review on the data quality and usage issues. *ISPRS Journal of Photogrammetry and Remote Sensing*.
- Mantelero, A. (2017). Personal data for decisional purposes in the age of analytics: From an individual to a collective dimension of data protection. *Computer Law & Security Review*, 32, 238-255.
- Marr, B. (2016, December 6). *What Is The Difference Between Artificial Intelligence And Machine Learning?* Retrieved from Forbes: <https://www.forbes.com/sites/bernardmarr/2016/12/06/what-is-the-difference-between-artificial-intelligence-and-machine-learning/>
- Matthews, K. (2018, October 11). *How GDPR is affecting big data ethics*. Retrieved from innovation enterprise : <https://channels.theinnovationenterprise.com/articles/how-gdpr-is-affecting-big-data-ethics>

- McClelland, C. (2017, December 4). *The Difference Between Artificial Intelligence, Machine Learning, and Deep Learning*. Retrieved from Medium: <https://medium.com/iotforall/the-difference-between-artificial-intelligence-machine-learning-and-deep-learning-3aa67bff5991>
- Metcalfe, J., Keller, E. F., & Boyd, D. (2016, May 23). *Perspectives on Big Data, Ethics, and Society*. Retrieved from Council for Big Data, Ethics, and Society: <https://bdes.datasociety.net/council-output/perspectives-on-big-data-ethics-and-society/>
- Miller, H. G., & Mork, P. (2012, November/December). From Data to Decisions: A value Chain for Big Data. *IT Professional*, 2-4.
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016, July-December). *Big Data & Society*, 1-21.
- Nunan, D., & Di Domenico, M. (2013). Market research and the ethics of big data. *International Journal of Market Research*, 55(4).
- O'Leary, D. E. (2013). Artificial Intelligence and Big Data. *IEEE Intelligent Systems*, 13, 1541-1672.
- O'Neil, C. (2016). *Weapons of Math Destruction*. Crown Books.
- Panijan, Z. (2010). Some Practical Experiences in Data Governance. *World Academy of Science, Engineering and Technology*, 62.
- Pediaa. (2016, July 17). *Difference Between Values and Principles*. Retrieved from Pediaa: <https://pediaa.com/difference-between-values-and-principles/>
- Phillips-Wren, G., & Jain, L. (2006). Artificial Intelligence for Decision Making Gloria. *Knowledge-Based Intelligent Information and Engineering Systems: 10th International Conference* (pp. 531-536). Bournemouth: Springer-Verlag.
- Quinn, M. J. (2004). *Ethics for the information age*. Pearson.
- Richards, N. M., & King, J. H. (2014). Big Data ethics. *Big Data & Society*, 49(2), 393-432.
- Rowley, J. (2012). Conducting research interviews. *Management Research Review*, 35(3/4), 260-271.
- Sagioglu, S., & Sinanc, D. (2013). Big Data: A Review. *2013 International Conference on Collaboration Technologies and Systems (CTS)*.
- Sessions, V., & Valtorta, M. (2006). The Effects of Data Quality on Machine Learning Algorithms. *Proceedings of the 11th International Conference on Information Quality, MIT*, (pp. 10-12). Cambridge MA, USA.
- Shadowen, A. N. (2017). Ethics and Bias in Machine Learning: A Technical Study of What Makes Us "Good". *CUNY Academic Works*.
- Siddaway, A. (2014). *What is a systematic literature review and how do I do one?*. Retrieved from Semantic Scholar: <https://www.semanticscholar.org/paper/WHAT-IS-A-SYSTEMATIC-LITERATURE-REVIEW-AND-HOW-DO-I-Siddaway/22142c9cb17b4baab118767e497c93806d741461>
- Singh, A. K., & Mishra, N. K. (2018). Ethical Theory & Business: A study based on Utilitarianism and Kantianism. *International Journal of Humanities and Social Development Research*, 2(1).
- Soiferman, L. K. (2010, April). *Compare and Contrast Inductive and Deductive Research Approaches*. Retrieved from ERIC: <https://eric.ed.gov/?id=ED542066>
- Stark, A. (1993, May-June). *What is the Matter with Business Ethics?* Retrieved from Harvard Business Review: <https://hbr.org/1993/05/whats-the-matter-with-business-ethics>
- Stoyanovich, J., Abiteboul, S., & Miklau, G. (2016). Data, Responsibly: Fairness, Neutrality and Transparency in Data Analysis. *International Conference on Extending Database Technology*. Bordeaux, France.

- Susser, D., Roessler, B., & Nissenbaum, H. (2019). Technology, Autonomy and Manipulation. *Internet Policy Review*, 8(2).
- Tavani, H. T. (2004). *Ethics and Technology*. Wiley.
- Tene, O., & Polonetsky, J. (2013, April). *Northwestern Journal of Technology and Intellectual Property*.
- Tiell, S., & O'Connor, L. (2016). *Building digital trust: The role of data ethics in the digital age*. Retrieved from Accenture: https://www.accenture.com/_acnmedia/pdf-22/accenture-data-ethics-pov-web.pdf
- UCL. (2017). *Anonymisation and Pseudonymisation*. Retrieved from UCL: <https://www.ucl.ac.uk/data-protection/guidance-staff-students-and-researchers/practical-data-protection-guidance-notices/anonymisation-and>
- Velasquez, M., Andre, C., Shanks, T., J., S., & Meyer, M. J. (2010, January 1). *What is Ethics?* . Retrieved from Markkula Center for Applied Ethics: <https://www.scu.edu/ethics/ethics-resources/ethical-decision-making/what-is-ethics/>
- Velasquez, M., Andre, C., Shanks, T., J., S., & Meyer, M. J. (2010, January 1). *What is Ethics?* Retrieved from Markkula Center for Applied Ethics at Santa Clara University: <https://www.scu.edu/ethics/ethics-resources/ethical-decision-making/what-is-ethics/>
- VPRO. (2018, April 22). *Slave to the algorithm*. Retrieved from VPRO: <https://www.vpro.nl/programmas/tegenlicht/kijk/Backlight/Slave-to-the-algorithm.html>
- Wahlstrom, K., Roddick, J. F., Sarre, R., Estivill-Castro, V., & deVries, D. (2006). On the Ethical and Legal Implications of Data Mining. *School of Informatics and Engineering Flinders University*.
- Ward, J. S., & Barker, A. (2013, September 20). *Undefined By Data: A Survey of Big Data Definition*. Retrieved from Cornell University: <https://arxiv.org/abs/1309.5821>
- Weinstein, B. (2017, October 31). *What's The Difference Between Ethics And Business Ethics?* Retrieved from Forbes: <https://www.forbes.com/sites/bruceweinstein/2017/10/31/whats-the-difference-between-ethics-and-business-ethics/#17097bf15428>
- Wells, D. (2018, May 2). *Data Ethics – The New Data Governance Challenge*. Retrieved from Erkerson Group: <https://www.eckerson.com/articles/data-ethics-the-new-data-governance-challenge>
- White, G., & Ariyachandra, T. (2016). Big Data and Ethics: Examining the Grey Areas of Big Data Analytics. *Issues in Information Systems*, 1-7.
- Whittlestone, J., Nyrupe, R., Alexandrova, A., & Cave, S. (2019). The Role and Limits of Principles in AI Ethics: Towards a Focus on Tensions. *The 2019 AAAI/ACM Conference*. Association for the Advancement of Artificial Intelligence.
- Wikipedia. (n.d.). *Data Governance*. Retrieved from Wikipedia: https://en.wikipedia.org/wiki/Data_governance#cite_note-3
- Wittmer, D. P. (2009). Behavioral Ethics in Business Organisations: What the Research Teaches Us. In J. O'Toole, & D. Mayer, *Good Business: Exercising Effective and Ethical Leadership* (pp. 62-73). Florence: Routledge.
- World Economic Forum. (2019, January). *AI Governance: A Holistic Approach to Implement Ethics into AI*. Retrieved from Weforum: <https://www.weforum.org/whitepapers/ai-governance-a-holistic-approach-to-implement-ethics-into-ai>
- Wright, D. (2011). A framework for the ethical impact assessment of information technology. *Ethics and Information Technology*(13), 199-226.
- Zwitter, A. (2014, July-December). Big Data ethics. *Big Data & Society*, pp. 1-6.

Annex A

Research Question 1				
Search Term 1		Search Term 2		Search Term 3
Big Data				
Big Data	AND	Ethics		
Big Data	AND	Analytics		
Responsible	AND	Data		
Data	AND	Ethics		
Big Data	AND	Privacy		
Big Data	AND	Privacy	AND	Ethics

Research Question 2				
Search Term 1		Search Term 2		Search Term 3
GDPR	AND	Regulation		
GDPR	AND	Principles		
Ethics				
Data	AND	Regulations		
Data	AND	Laws		
Ethical	AND	Culture		
Big Data	AND	Politics		
Big Data	AND	Ethics		
Big Data	AND	Ethics	AND	GDPR
Data	AND	Governance		
Big Data	AND	Governance		
Ethics	AND	Regulations	AND	Data
Ethics	AND	Regulations		
Ethics	AND	Law		

Research Question 3				
Search Term 1		Search Term 2		Search Term 3
Big Data	AND	Ethics		
Big Data	AND	Analytics	AND	Ethics
Big Data	AND	Challenges		
Big Data	AND	Privacy		
Big Data	AND	Bias		
Data	AND	Governance	AND	Ethics
Responsible	AND	Data	AND	Analysis
Data	AND	Management	AND	Ethics
Digital	AND	Trust		
Artificial Intelligence	AND	Bias		
Artificial Intelligence	AND	Ethics		
Algorithm	AND	Bias		
Computer	AND	Ethics		
Machine Learning	AND	Ethics		
Data Mining	AND	Ethics		
Data Mining	AND	Privacy	AND	Threats
Data Mining	AND	Discrimination		

Annex B

Semi-Structured Interview Protocol
1) Introduction
<p><i>Permission to record the interview is asked to the interviewee, with the motivation that a transcription of the interview will allow a more accurate analysis of the results.</i></p> <ul style="list-style-type: none"> • Can you introduce yourself?
2) Definition of Big Data Ethics
<p><i>The topic of the research is introduced: the researcher is investigating the ethics of Big Data and wants to identify a solution to unethical Big Data activities in the form of Data Governance practices.</i></p> <ul style="list-style-type: none"> • Did you ever deal with the ethics of data in your work? If so, can you give me an example of a project? <p><i>The proposed definition of Big Data Ethics from the literature review is introduced.</i></p> <ul style="list-style-type: none"> • Do you think that it is important for an organisation to address the dimension of ethics when using Big Data for decision making purposes?
3) Ethical principles applied to Big Data
<p><i>It is explained how in the literature review part of the research fundamental ethical principles have been applied to the context of Big Data to determine unethical Big Data activities.</i></p> <p><i>A printed version of the literature review framework is shared. The interviewee is shown and explained the Big Data lifecycle and asked to think of the presented Big Data activities in terms of the Big Data lifecycle.</i></p> <p><i>The definition of 'Insert_principle' proposed in literature is introduced, together with the sub-principle related to it and the consequent unethical Big Data activity.</i></p> <ul style="list-style-type: none"> • Do you see this activity potentially violating the related sub-principle? • Do you have anything to add regarding this? <p><i>The questions above are asked for all other sub-principles.</i></p> <ul style="list-style-type: none"> • Do you find that these sub-values share a connection with the principle 'insert_principle'? Is there anything you would change or you think is missing? <p><i>The questions of section 3 are repeated until all principles are discussed.</i></p>
4) Critical Big Data activities
<ul style="list-style-type: none"> • Given the activities that we discussed so far, if you had to appoint 4 or 5 critical ones that you as a company would want to address first, what would those be? <i>Follow-up: Given the Big Data lifecycle presented earlier in Annex B, which phases would you suspect to be the most high-risk from an ethical perspective?</i> • If you were responsible for addressing these unethical activities within a company, how would you do it?
5) Feedback
<ul style="list-style-type: none"> • Do you have any feedback on the questions and the way the interview was conducted?

Annex C

Codes		+
Role of laws in the ethical debate	3	
Importance of addressing data ethics	9	
Feedback on Autonomy	0	
Agree with relationship	25	
Disagree with relationship	3	
General recommendations	8	
Feedback on Nonmaleficence	0	
Agree with relationship	13	
Disagree with relationship	1	
General recommendations	7	
Feedback on Beneficence	0	
Agree with relationship	13	
Disagree with relationship	1	
General recommendations	11	
Feedback on Justice	0	
Agree with relationship	7	
Disagree with relationship	0	
General recommendations	3	
Feedback on Privacy and Data Protection	0	
Agree with relationship	31	
Disagree with relationship	10	
General recommendations	16	
General feedback on framework	8	

Annex D

Yes, I do. Informed Consent in itself is such a difficult topic because what is really informed and what is consent?

We all have the example of terms and conditions that we all don't read and give consent to. This is an important topic for Respect for Autonomy and the relationship that you found is there, in my opinion.

Interviewee 8 agreed with the example of the Informed Consent category.

"About 10 times a day I get a message from my browser that some website would like to send me messages: this shouldn't really be allowed." (Transcript 8 Complete, Paragraph 32)

Transcript 8 Complete, Paragraph 30

when I first read it I thought that this is closely related to Autonomy, because I want to know what is available for me as a person, and make my own decisions. Autonomy and the Inclusion categories would be me breaking out of my Facebook bubble, my Spotify bubble, and really being autonomous in that digital realm: making my own choices and having all the available information handed to me.

Interviewee 8 disagreed with the example of the Social Solidarity, Inclusion and Exclusion bubble.

"A person should have the Autonomy to say that they don't want to use Facebook, and still not miss out on anything" (Transcript 8 Complete, Paragraph 35) "and not only be provided with the information that was selected for their profile. I want my profile to be inclusive. If I only listen to hip hop music, I would still love for Spotify to show me the charts of what people are listening to, in which there is also EDM music, rock music and not only hip hop." (Transcript 8 Complete, Paragraph 36)

The recommendation is to revise the elaboration of the principle.

Annex E

Data Governance Practices for Big Data ethics – Online Questionnaire	
Introduction to Survey	
<p>Dear Respondent,</p> <p>The survey is part of a Master thesis by the name of “Data Governance practices for Big Data ethics”, as part of the ‘ICT in Business’ course in Leiden University. The researcher is studying ethical problems in the business environment raised by the use of Big Data for decision making purposes.</p> <p>This survey aims to investigate the major data-related ethical issues in commercial companies. Lastly, this survey has the goal of collecting information regarding the governance practices currently in place within commercial companies to address those ethical issues.</p> <p>The survey should take 20-25 minutes to complete. Be assured that the results will be kept in the strictest confidentiality and any personal information or result concerning the ethical evaluation of your company will be kept anonymous in the final thesis document.</p> <p>Thank you for your participation and for investing your time in this research!</p> <p>DISCLAIMER: This research is being conducted with the help of the Enterprise Architecture team within Deloitte. The survey is however not being sent on behalf of Deloitte and the results will strictly be used by the researcher for the completion of their Master thesis.</p>	
Background Questions	
<i>The following questions will allow the respondent to provide information regarding their background and the context they work in.</i>	
In what industry is the company you work for active in?	<ul style="list-style-type: none"> • Financial services • Telecommunication • Transportation • Retail • Pharmaceutical • Construction • Energy • IT • Other (please specify)
How many employees does your company have?	<ul style="list-style-type: none"> • Less than 10 • Between 10 and 50 • Between 50 and 250 • More than 250
What is your role within your company?	
How many years of working experience do you have?	<ul style="list-style-type: none"> • Less than 5 years • Between 5 and 10 years • Between 10 and 20 years • More than 20 years

Ethical Evaluation

Commercial organizations often put a lot of effort into developing their analytics capabilities, yet are struggling to use them in an ethical matter. In this research, the concept of **Big Data ethics** has been introduced as a discipline that studies ethical problems that arise in the business environment when using Big Data and algorithms to make decisions; its goal is to develop moral rules, standards or practices that support moral decision making based on Big Data analytics - so that ethical principles, human rights and data regulations are not violated throughout the data lifecycle.

Down below is the data lifecycle description which has been used throughout the research, which will also be used as a reference in this survey.



The following sections aim to investigate the respondents' perceived ethical problems in using Big Data within commercial organizations to support decision making, as well as the practices they might use to address them, throughout the phases of the data lifecycle.

PLEASE NOTE that the questions in the survey have been constructed on the basis of previously executed interviews with experts and are not necessarily consistent in the number of possible answers and/or the elaboration of the answers. Furthermore, not all phases will be addressed in this survey, but only the ones that were deemed to be higher-risk from an ethical perspective by previously interviewed experts.

Data Collection phase

The data collection phase consists in collecting raw data from all possible - and relevant - sources. This phase is also focused on turning the collected unstructured data into structured data.

Did you ever encounter ethical problems or dilemmas associated with the data collection phase in your company?

- Yes
- No

If previous answer is Yes: Which ethical problems or dilemmas associated with the data collection phase did you encounter?

Do you have any practices in place within your company to address and deal with the ethical problems associated to the data collection phase?

- Yes
- Maybe
- No

If previous answer is Yes or Maybe: Which practices does your company have in place to address ethical problems or dilemmas associated with the data collection phase? (please select max. 3 options)

- Define a data strategy
- Reduce error of data by looking for biases in the way data is collected
- Openness towards the customers and authorities
- Accountability for the data collection
- Establish effective policies and procedures to guarantee alignment between business principles and the collection of data
- Design a review board with the power to approve or deny the collection of new data
- Accompany data with information about its provenance and processing
- Evaluate data for fidelity to the phenomenon under consideration
- Other _____

<p><i>If previous question was shown:</i> Can you elaborate on what additional or alternative measure you would take to collect data responsibly and ethically, which hasn't been mentioned in the previous question?</p>	<p>_____</p> <p>_____</p> <p>_____</p>
<p>Data Cleansing phase The cleansing of data consists in detecting and correcting corrupt or inaccurate records from a data, which saves calculation time and space.</p>	
<p>Did you ever encounter ethical problems or dilemmas associated with the data cleansing phase in your company?</p>	<ul style="list-style-type: none"> • Yes • No
<p><i>If previous answer is Yes:</i> Which ethical problems or dilemmas associated with the data cleansing phase did you encounter?</p>	<p>_____</p> <p>_____</p>
<p>Do you have any practices in place within your company to address and deal with the ethical problems associated to the data cleansing phase?</p>	<ul style="list-style-type: none"> • Yes • Maybe • No
<p><i>If previous answer is Yes or Maybe:</i> Which practices does your company have in place to address ethical problems or dilemmas associated with the data cleansing phase? (please select max. 3 options)</p>	<ul style="list-style-type: none"> • Establish effective policies and procedures to guarantee alignment between business principles and the data cleansing phase • Validate assumptions baked into the normalization methodology • Define a data strategy to clarify how the company cleanses data • Look for systematic biases in the way data is cleansed and validate data cleansing assumptions • Other _____
<p><i>If previous question was shown:</i> Can you elaborate on what additional or alternative measure you would take to cleanse data responsibly and ethically, which hasn't been mentioned in the previous question?</p>	<p>_____</p> <p>_____</p> <p>_____</p>
<p>Data Analysis phase The data analysis phase consists in analyzing raw data to draw information and knowledge from it. One of the most important steps of data analysis is the selection of appropriate techniques for data analysis (which can be Artificial Intelligence/Machine Learning or another methods).</p>	
<p>Did you ever encounter ethical problems or dilemmas associated with the data analysis phase in your company?</p>	<ul style="list-style-type: none"> • Yes • No
<p><i>If previous answer is yes:</i> Which ethical problems or dilemmas associated with the data analysis phase did you encounter?</p>	<p>_____</p> <p>_____</p>
<p>Do you have any practices in place within your company to address and deal with the ethical problems associated to the data analysis phase?</p>	<ul style="list-style-type: none"> • Yes • Maybe • No
<p><i>If previous answer is Yes or Maybe:</i> Which practices does your company have in place to address ethical problems or dilemmas associated with the data analysis phase? (please select max. 3 options)</p>	<ul style="list-style-type: none"> • Reducing error of analysis • Accountability of data usage • Introduce standardized operations to evaluate algorithms • Establish a common understanding of algorithms • Establish effective policies and procedures to guarantee alignment between business principles and the data analysis phase

	<ul style="list-style-type: none"> • Define a data strategy to clarify how the company works with data • Designate a cross-functional review board with the power to approve or deny the deployment of insights from the analysis of sensitive questions • Look for systematic biases in the data analysis phase • Other _____
<i>If previous question was shown:</i> Can you elaborate on what additional or alternative measure you would take to analyses data responsibly and ethically, which hasn't been mentioned in the previous question?	<p>_____</p> <p>_____</p> <p>_____</p>
<p>Data Visualization phase</p> <p>The data visualization phase consists in displaying the analytic results in a clever way and present them to decision makers. The goal is to turn meaningful information in a format that decision makers can easily understand and consume to make decisions.</p>	
Did you ever encounter ethical problems or dilemmas associated with the data visualisation phase in your company?	<ul style="list-style-type: none"> • Yes • No
<i>If previous answer is Yes:</i> Which ethical problems or dilemmas associated with the data visualization phase did you encounter?	<p>_____</p> <p>_____</p>
Do you have any practices in place within your company to address and deal with the ethical problems associated to the data visualization phase?	<ul style="list-style-type: none"> • Yes • Maybe • No
<i>If previous answer is Yes or Maybe:</i> Which practices does your company have in place to address ethical problems or dilemmas associated with the data visualization phase? (please select max. 3 options)	<ul style="list-style-type: none"> • Awareness of human bias which can affect the way results are interpreted • Other _____
<i>If previous question was shown:</i> Can you elaborate on what additional or alternative measure you would take to analyses data responsibly and ethically, which hasn't been mentioned in the previous question?	<p>_____</p> <p>_____</p> <p>_____</p>
<p>Decision Making phase</p> <p>The decision making phase consists in determining what actions or decisions can be made on the basis of the analyzed data and visualized results.</p>	
Did you ever encounter ethical problems or dilemmas associated with the decision making phase in your company?	<ul style="list-style-type: none"> • Yes • No
<i>If previous answer is Yes:</i> Which ethical problems or dilemmas associated with the decision making phase did you encounter?	<p>_____</p> <p>_____</p>
Do you have any practices in place within your company to address and deal with the ethical problems associated to the decision making phase?	<ul style="list-style-type: none"> • Yes • Maybe • No
<i>If previous answer is Yes or Maybe:</i> Which practices does your company have in place to address ethical problems or dilemmas associated with the decision making phase? (please select max. 3 options)	<ul style="list-style-type: none"> • Establish effective policies and procedures to guarantee alignment between business principles and decision making phase • Establish a common understanding of how specific decisions are made • Have data scientists validate predictions continually

	<ul style="list-style-type: none"> • Monitor the performance of a system after lunch by means of black box • Test for feedback loops and the possibility that mistaken decisions might disproportionately harm individuals • Designate an internal role responsible for owning the outcomes of analysis • Other _____
<i>If previous question was shown:</i> Can you elaborate on what additional or alternative measure you would take to make decisions responsibly and ethically, which hasn't been mentioned in the previous question?	<hr/> <hr/> <hr/>
<p>Data Archival and Data Destruction phases</p> <p>The data archival and data destruction phases respectively consist in archiving and disposing of the data after its exploitation.</p>	
Did you ever encounter ethical problems or dilemmas associated with the data archival and data destruction phases in your company?	<ul style="list-style-type: none"> • Yes • No
<i>If previous answer is yes:</i> Which ethical problems or dilemmas associated with the data archival and data destruction phases did you encounter?	<hr/> <hr/>
Do you have any practices in place within your company to address and deal with the ethical problems associated to the data archival and data destruction phases?	<ul style="list-style-type: none"> • Yes • Maybe • No
<i>If previous answer is Yes or Maybe:</i> Which practices does your company have in place to address ethical problems or dilemmas associated with the data archival and data destruction phases? (please select max. 3 options)	<ul style="list-style-type: none"> • Consider the risk that retained data could be reidentified • Understand how and why data must be retained and know how it will be used, to know when it can be profitably discarded or properly anonymized • Other _____
<i>If previous question was shown:</i> Can you elaborate on what additional or alternative measure you would take to archive and destruct data responsibly and ethically, which hasn't been mentioned in the previous question?	<hr/> <hr/> <hr/>
<p style="text-align: center;">Contextual Questions</p> <p>The following questions aim to evaluate the respondent's awareness of the ethical concerns and risks raised by using Big Data for decision making purposes within commercial organizations, as well as the role of both organizations and governmental institutions in addressing the ethical behavior from a data perspective.</p>	
How risky do you think it is for your company to not address the ethical problems deriving from using Big Data for decision making purposes?	<ul style="list-style-type: none"> • Extremely risky • Very risky • Moderately risky • Slightly risky • Not risky at all
<i>If previous answer is Extremely risky, Very risky, Moderately Risky or Slightly risky:</i> What do you think are the risks your company faces by handling data unethically? <i>If previous answer is Not risky at all:</i> What do you think are the risks your company faces by handling data unethically?	<hr/> <hr/> <hr/> <hr/>

How would you rank your perceived importance of each data lifecycle phase in addressing the ethical problems or dilemmas deriving from the use of Big Data (please rate each individual phase from 1. phase of highest importance to 6. phase of lowest importance)? Please note: each phase needs to be rated with a different level of importance)	Data collection (1 to 6) Data cleansing (1 to 6) Data analysis (1 to 6) Data visualization (1 to 6) Decision making (1 to 6) Data archival and data deletion (1 to 6)
How would you rank the priority level within your company of each data lifecycle phase in addressing the ethical problems or dilemmas deriving from the use of Big Data (please rate each individual phase from 1. phase of highest priority to 6. phase of lowest priority)? Please note: each phase needs to be rated with a different level of priority)	Data collection (1 to 6) Data cleansing (1 to 6) Data analysis (1 to 6) Data visualization (1 to 6) Decision making (1 to 6) Data archival and data deletion (1 to 6)
To what extent do you think that laws and regulations are responsible in establishing ethical behavior within an organization from a data perspective?	<ul style="list-style-type: none"> • Fully responsible • Mostly responsible • Moderately responsible • Slightly responsible • Not responsible at all
To what extent do you think that it is a company's own responsibility to establish ethical behavior from a data perspective?	<ul style="list-style-type: none"> • Fully responsible • Mostly responsible • Moderately responsible • Slightly responsible • Not responsible at all
Feedback	
Thank you for completing the survey! If you have any comments on the survey, please leave them down below. If you wish to receive the results of the research, feel free to contact me at bevacqua.nunzia@gmail.com	
Do you have any feedback on the survey?	_____

Annex F

Requirement 1: Job Title
<ul style="list-style-type: none"> • Chief data officer • Data governance officer/head/director/manager/specialist • Data protection/privacy officer/head/director/manager/specialist • Data ethics officer/head/director/manager/advisor/specialist • Ethics and compliance officer/head/director/manager/specialist • Compliance and ethics officer/head/director/manager/specialist • Data scientist (senior and non) • Data analyst (senior and non) • Data architect • Data engineer
Requirement 2: Industry
<ul style="list-style-type: none"> • Financial services, banking and insurance (e.g. ABN Amro, Rabobank, ING, Nationale-Nederlanden, Aegon, Allianz) • Telecom (e.g. Vodafone, KPN, Liberty Global) • Transportation (e.g. FedEx, PostNL) • Retail (e.g. Nike, Asics, bol.com, HEMA, Philips, Coca Cola, IKEA, Jumbo, Albert Heijn) • Pharma (e.g. Pfizer, Johnson & Johnson)

Note: the examples listed in this section are only an indicative guidelines of companies which are likely to collect and use personal data to do analysis and generate decisions for the business. They do not mean to correspond to the actual companies surveyed in this study.

Requirement 3: Field of expertise

- GDPR compliance
- Data Governance
- Data ethics
- Digital ethics
- Data privacy and protection
- Data science
- Data analysis

Annex G

Object: Leiden Student – Master thesis survey request – Ethics of Big Data

Dear First_Name Last_Name,

I am Nunzia Bevacqua, a student enrolled in the Master course ICT in Business in Leiden University. I am currently writing my Master thesis on the ethics of Big Data and how to address them by means of Data Governance practices.

I am conducting a survey in which I want to explore and validate known ethical problems concerning the use of Big Data for decision making, as well as investigate how such ethical problems are being addressed within organizations.

I am looking for a candidate that would be willing to fill out my survey and I believe you would be a perfect choice! I expect the duration of the survey of being around 20 minutes, and the results will be kept confidential. If you were agree to help me out, I would like to receive your response before February 7th. Furthermore, if you know anybody in your network that deals with data ethics, data privacy and/or data protection, could you share it with them as well? Your help will be fundamental in reaching out my target number of respondents.

This is the link to the survey: https://leidenuniv.eu.qualtrics.com/jfe/form/SV_1GhWtmwJMzBTFzf I would be grateful if you could take the time to help me out. To express my gratitude, I can offer to share the results of the thesis with you once it is completed.

If you have any doubts or questions, do not hesitate to contact me!

Kind regards,

Nunzia Bevacqua

Annex H

Semi-Structured Validation Protocol																						
1) Introduction																						
<p>Permission to record the interview is asked to the validation panel, with the motivation that a transcription of the session will allow a more accurate analysis of the results.</p> <p>The research topic is introduced, and the structure of the Data Governance framework for Big Data ethics is explained.</p>																						
2) Framework Validation																						
<p>The panel is shown an overview of the identified activities and practices for each phase of the Big Data lifecycle.</p> <p>The panel is shown each addition to the framework that requires validation for a specific phase of the Big Data lifecycle. The following questions are asked:</p> <ul style="list-style-type: none"> • Is the highlighted practice valid from a Data Governance point of view? • Is the highlighted activity valid from an ethical point of view? • Does it help address the corresponding Big Data activity? <p>The panel is shown a list of practices used within the surveyed companies, from the most used to the least used, referring to a specific phase of the Big Data lifecycle. The following questions are asked:</p> <ul style="list-style-type: none"> • Can you explain why the first x practices are the most used? • Are you surprised by any of the results (e.g. is there any practice you would have expected to be more or less used)? 																						
3) Big Data lifecycle priorities																						
<p>The panel is engaged into an exercise to establish a set of criteria to rank the lifecycle phases. This will support the explaining of how the survey respondents came up with their answers.</p> <p>Please think of a set of criteria you would use to rank the Big Data lifecycle phases from the most risky (and therefore the most important to address) to the least risky (and therefore the least important to address) for commercial companies.</p> <p>Please use the criteria you thought of to rank the Big Data lifecycle phases from the most risky to the least risky from the perspective of commercial companies.</p>																						
																						
<p>The panel is shown the opinions of the experts compared to the opinions of the survey respondents regarding the priorities of the lifecycle phases. This result is shown to determine where companies currently stand in addressing Big Data ethics.</p>																						
	<table border="1"> <thead> <tr> <th></th> <th>Level of priority of lifecycle phases within surveyed organisations</th> <th>Riskiness level of lifecycle phases (expert interviews)</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>Data visualization</td> <td>Data collection</td> </tr> <tr> <td>2</td> <td>Data archival and data deletion</td> <td>Data analysis</td> </tr> <tr> <td>3</td> <td>Data cleansing</td> <td>Decision making</td> </tr> <tr> <td>4</td> <td>Data analysis</td> <td>Data cleansing</td> </tr> <tr> <td>5</td> <td>Data collection</td> <td>Data visualization</td> </tr> <tr> <td>6</td> <td>Decision making</td> <td>Data archival and data deletion</td> </tr> </tbody> </table>		Level of priority of lifecycle phases within surveyed organisations	Riskiness level of lifecycle phases (expert interviews)	1	Data visualization	Data collection	2	Data archival and data deletion	Data analysis	3	Data cleansing	Decision making	4	Data analysis	Data cleansing	5	Data collection	Data visualization	6	Decision making	Data archival and data deletion
	Level of priority of lifecycle phases within surveyed organisations	Riskiness level of lifecycle phases (expert interviews)																				
1	Data visualization	Data collection																				
2	Data archival and data deletion	Data analysis																				
3	Data cleansing	Decision making																				
4	Data analysis	Data cleansing																				
5	Data collection	Data visualization																				
6	Decision making	Data archival and data deletion																				

- What can we say about the status of the surveyed companies in addressing Big Data ethics?
- Did the respondents use different criteria to make the judgement? What could this criteria be?
- Why are their priorities not aligned with the ideal scenario?

The panel is shown the opinions of the experts compared to the opinions of the survey respondents regarding their perceived importance of the lifecycle phases. This result is shown to determine where companies might be headed in their ethical journey.

	Perceived importance of lifecycle phases by survey respondents	Riskiness level of lifecycle phases (expert interviews)
1	Data cleansing	Data collection
2	Data visualization	Data analysis
3	Data archival and data deletion	Decision making
4	Data collection	Data cleansing
5	Data analysis	Data visualization
6	Decision making	Data archival and data deletion

- What can we say about the discrepancy between the expert opinions and the survey respondents?
- Why are the results not aligned?

4) Evaluation of framework usability

- Can the framework help organizations assess how they deal with Big Data ethics?
- Who would be a good candidate to use the framework (e.g. role, company)?

Annex I

Respect for Autonomy							
Autonomy		Dignity		Informed Consent		Social Solidarity, Inclusion and Exclusion	
A/D	Additional comments	A/D	Additional comments	A/D	Additional comments	A/D	Additional comments
Agree	Having personalization algorithms in place causes people to only see the things a third party thinks are good for them: users' decisions are pushed in a certain direction by the third party, which imposes on them what they think they must like and consequent decisions of theirs such as a purchasing act. The algorithm is in a way making a choice for the user.	Agree	Big Data adds a dimension of complexity to the act of personalization: right now much more data is available regarding people which allows to do the same but on a bigger scale.	Agree	It costs the user too much time to look into the consent documents. The user wanted to directly access the service and policy notices constitute a barrier from what the user intends to do on a certain website: this causes the user to click it away as fast as possible without reading it first.	Agree	Additional examples that motivate representativeness problems in the datasets: 1) People with a certain interest are probably going to visit certain sites, which causes those sites to only collect data and opinions about selected groups of people. 2) People without a job spend more time on social media than those who work all day, which causes these sites to gather more information about unemployed people than working people.
Agree	Currently users give away their data for free to companies. People should have some kind of data personality that has rights and says whether or not a certain company can or cannot use their data.	Agree	Additional example: if a person doesn't have enough money to live a healthy life then they will probably be less healthy and insurance companies might charge them more money.	Agree	The problem does not only concern whether consent is asked for, but also the way it is done. Even when companies ask for informed consent to the user there may be barriers in the way they ask for it (for example by forcing the user to read long and complicated privacy policies). That behaviour is rather unethical.	Agree	Users cannot get the right offers if they are not included in the data sets.
Agree	In the past marketing was 1 to many: it was broadcasted and it was never fully matched to one's personality. With Big Data advertising has become personalized, and it has become so smart that A: the user does not perceive that they are being influenced; and B: they cannot resist it because the Big Data system knows so much about their personality that it is able to manipulate them.	Agree	What hurts the dignity of a person is feeling like they are being treated like the outcome of a model: this feels degrading for them and at worse if the model is black and white (e.g. good or bad, accepted or not accepted) the user is being reduced to a mere number.	Disagree	The other categories in the framework are clear values, but this category is more of a legal term. This category sounds like a power imbalance, which means that the data collector leaves the user in the dark, does whatever they want with the user's data and may inform them with privacy notices but the user might not have a clue about what the notice is really saying.	Agree	Additional examples: 1) This is particularly a problem with face recognition technologies: many of the data sets are collected from white people and it is well known that black people are less recognized by face recognition. 2) Handicapped people or people affected by mental illness are not recognized at all: it is like they don't exist because they are not in the data.

Agree	Whenever information arrives to a user, that information will in a sense manipulate them because they will form a stance towards it, whether it is positive or negative. However, when companies are deliberately feeding the users a certain kind of information to move them in a certain direction, such behavior limits the users' autonomy.	Agree	The categorization of people and act of putting them into certain buckets can limit their ability to do certain things, and there is a danger in this type of situation. What makes the situation worse is when a person is in a certain category due to factors they cannot influence e.g. the area where they live where perhaps other people behave in a certain way: the system might think that their behavior is likely to match the one of the people they live nearby to.	Agree	People are not rational: when they have to agree that their data is being collected by a website, even when they give consent to it they don't understand the implications of that action and what their information will be used for (e.g. the fact that such data is used to make a profile of them across several other websites).	Agree	The possibility of unethical usage is always there: when people are not filtered out of the data intentionally (but instead it just so happens that certain people are not in the dataset), that behavior would not be unethical. If the exclusion of individuals is done intentionally, that behavior would on the other hand be unethical.
Agree	The GDPR does not allow automated decision making including profiling. If the use of an algorithm deals to a decision then the user has the right to human intervention (for example if a person is only allowed to receive social security based on the decision of an algorithm, which could be harmful to their personal situation).	Agree	Due to the GDPR (Article 9) companies have to comply to extra criteria when processing sensitive personal data (such as age, gender, racial background and disabilities).	Agree	According to Article 4 of the GDPR "consent must be given freely, it must be specific, informed and unambiguous": if a user has to read a complicated 100 pages document to be informed, that would be unlawful. Consent is one of the key principles of the GDPR, but it is only one of the possible six lawful basis stated by the GDPR that allow a company to collect personal data.	Agree	/
Agree	/	Disagree	No ethical principle is being violated through price differentiation based on data: using the same line of thought also the differentiation of tax rates based on income would be a violation of dignity.	Agree	Despite the fact that consent should be informed, in reality every website has a cookie notice which does not provide informed consent. The GDPR created a paradox by coming up with very strong demands, which now causes people to give consent to everything due to the invasive cookie notices: this ultimately does not allow the informed consent principle to work correctly anymore.	Agree	Additional example: Due to banks closing in smaller villages to promote the use of online banking, older people that don't have a mobile phone cannot do their banking anymore. The fact that some people do not use certain services might cause problems with Big Data systems.
/	/	/	/	/	/	/	/

Agree	/	Agree	/	Agree	It is hard to determine what is truly informed consent. The interviewer, as a user, gets messages from his browser about 10 times a day that some website would like to collect information about him or send him messages: such behavior should not be allowed.	Disagree	Inclusion is a concept close to Autonomy: it means wanting to know what is available for me as a person, and make my own decisions. Inclusion is the empowerment of the user to break out of their Facebook bubble, or Spotify bubble, and be autonomous in that digital realm: it means being able to make your own choices as a user and having all the available information handed to you. A person should have the autonomy to say that they don't want to use Facebook, and still not miss out on anything and not only be provided with the information that was selected for their profile.
-------	---	-------	---	-------	--	----------	---

Nonmaleficence			
Safety		Discrimination and social sorting	
A/D	Additional comments	A/D	Additional comments
Agree	By putting people in a certain corner based on characteristics such as location or income, companies may put their safety at risk.	Agree	Companies should put algorithms in place that don't take discriminating characteristics into consideration: they should be excluded from the analysis to not violate this ethical principle. However, the context can also be important to determine which personal characteristics to consider into the analysis and which not: if the purpose of the analysis will genuinely benefit the person, then it should be allowed to use certain characteristics: if the purpose is, for example, to only send advertisements then it should not be allowed.
Agree	Additional example: self-driving cars is an example of programmed ethics, and it is about what you want the AI system to decide. It is however not related to decision making based on data.	Agree	Additional example: by letting Artificial Intelligence decide who is to lead a company based on historical data, then a good CEO would probably be a male.
Agree	Companies overlook the fact that models are flawed: every model has instances in which they do not predict well. Companies should ask what the impact is of a false positive on a person (with the potential safety risks on them), and how they are going to protect people from the flaws of the model.	Disagree	Some distinctions are okay to make, while some others are not. There are some differences generally accepted by the public, such as the fact that a 40 year old driver is a far safer risk on average than a 20 year old driver. It is crucial, however, that distinctions are not made based on characteristics that people can't influence.
Agree	If algorithms don't work well, they put the safety of people in danger. Additional example: there are situations in which people get a certain medical treatment because a Big Data system predicts that it will be likely to succeed on them, therefore the safety of people is in these cases put into the hands of algorithms.	Agree	It is totally unknown how Big Data systems work, whether they work correctly or if they are of any actual use (e.g. systems used by the government that decide whether people are allowed to fly or not): these systems might be sometimes biased and discriminate people.
Agree	/	Agree	/
Agree	Once a Big Data-based decision has been made, there is no check happening to determine whether the decision generated by the analysis of a Big Data set is true or not. False positives can ultimately affect the safety of people. Proposed an alternative example.	Agree	Placing people into categories can work to their advantage or disadvantage: if people are put into a position of disadvantage, there is a violation of ethics. Additional example: insurance companies were using a model that asked for the gender of people and generated a more expensive car insurance for women, because women were deemed to be worse drivers by the system.

Agree	/	Agree	/
Agree	/	Agree	Additional example: the police in Rotterdam uses a profiling system to profile high risk areas within the city; this profiling was apparently discovered to profile in the wrong way, with consequences on the people.

Beneficence			
Imbalance of Power		Value Sensitive Design	
A/D	Additional comments	A/D	Additional comments
Agree	Users give data to companies without getting anything back.	Agree	There is the need of having a human involved in the process checking on the algorithm, because allowing the algorithm to take control of the process does not ensure that the results will be good 100% of the time. The human component is important to give direction to the algorithm and prevent algorithm mistakes.
Agree	Companies hold the informational power over the data subjects. The whole business model for a lot of Big Data companies is "You give me something and in return I get something back that you don't even realize the value of". If companies use people strictly for their data and for themselves in their own interest, then they violate this ethical principle.	Agree	The integration of values also depends on the moment the person made the selection of the information (factors such as whether the person is hungry or not, whether it's summer or winter, whether they have slept well or not) and on what comes to mind at that specific moment. These factors influence the way values are integrated into a system.
Agree	Beneficence means doing well and organizations should identify and implement possibilities to do well with data. Ethical practices should be promoted, which means that organizations should use customer data not only for their own good but also purely for the customer. For example, if a company were to think in a care ethics type of way they will not put privacy notices that nobody can understand. However, the current situation is one of imbalance where the user only gives data away, the organization takes it and does whatever it wants with it; the user doesn't see anything back from it and hopes he won't be harmed too much by the consequences.	Agree	The shaping of the design from the human is almost inevitable: when making something they will, consciously or not, put some of their values into it. Accountability is the main counterbalance for this: if you can explain and will explain what you have done you are open for debate and anyone can say whether they agree or not; if you don't have accountability you are stuck with the values of the person that designed the system in the first place.
Agree	Companies gather a lot of data but they might only use it to send out promotions to the user; as a user, once you have given out your data it is out of your hands and you cannot control anymore what you receive because of it.	Agree	The goal of most algorithms is that they work as well as possible: if errors are introduced in them and people are harmed as consequence of it, that would be unethical. However, a situation in which people have a positive goal in mind and something goes wrong is different than when people have bad intentions and of course things will go wrong.
Agree	The fact that big tech companies own huge amounts of personal data creates an imbalance of power. However, at the same time if the law works the way it should work, the situation would be more balanced. The law should have the responsibility to create a situation of balance: the purpose of the GDPR for example is to give users control over their own data, which could lead to a more balanced situation.	Agree	If the input is discriminatory it will also give a discriminatory result.
Disagree	If there wasn't any reciprocity going on there would be an imbalance of power. However, there is a reciprocity going on when services like Google take the personal data of users and give back a free email service, a free navigation service etc. which are of huge benefit to society.	Agree	Designing something with one's own perspective (such as thinking it might be handy to include certain information in the analysis) which only takes their interests into consideration, is unethical in that it does not consider the interests of the user.

Disagree	In every power relation there is always an imbalance of power: wanting to avoid all imbalances of power in the sense of things always being equal would mean abolishing all power relationships, so that nobody has power over nobody. There will always be a difference of power between the government and the citizen, but the citizen also has ways to counteract that power in order to avoid an abuse of it. There is a growing imbalance of power between companies and the government: big tech companies abused the fact they had the information power to gain even more power for themselves.	Agree	A big part of value sensitive design is about finding out what are the values that a product needs to fulfill for all stakeholders and consciously putting those things in the design, in a beneficence kind of way. People have to realize that when they design they put a normative stance into that design: however value sensitive design is about consciously designing to achieve certain values for the stakeholders.
Agree	/	Agree	/

Justice

Equality and Fairness

A/D	Additional comments
Agree	Big Data is making personalization easier and increases the scale of this phenomenon. This phenomenon is unfair and legislation should be introduced to make this behavior illegal. <i>Additional example:</i> Albert Heijn sends customers personalized discounts which means that I might have to pay a different amount than somebody else based on the products that I bought before. A similar phenomenon happens with airline tickets and vacations.
Agree	It is possible to look at the behavior of using personalization algorithms from the point of view of fairness. However, it is not the responsibility of the business to create situations of fairness, but it is rather a political matter. Saying that a behavior is unfair could lead to a law that forbids such behavior. In the market you could create a situation of fairness by having new participants that target the groups discriminated by other companies.
Agree	Justice is about what society expects, which also changes in time: if people expect that they pay the same price of a plane ticket as someone else, but actually they don't and find out, then they will feel cheated; if on the other end they expect it, then it is fine. Furthermore, everybody has some idea of what these expectations are: for example, we should treat women the same as men, or we should not fool people that think they are getting a fair price by making them pay double the price of their neighbor.
Agree	It is not fair that if a system thinks I can spend more money on a ticket, I should receive a higher price for it. It is arguable in the sense that the user decides for himself what they want to pay and the company decides for itself what it wants to ask as a price: but this impacts the equality of people.
Agree	It would be unfair if, for example, a woman were to buy a man's deodorant and it would be cheap for them because they are not supposed to be interested in it, whereas for a man the same product would be more expensive. This type of behaviour might cause harm or disadvantage.
Agree	<i>Additional example:</i> it is unfair if one goes to the Albert Heijn and gets a certain pricing presented because they have information about their previous sales activity, but their neighbor paid less for the same product.
Agree	/
Agree	Hardcore justice such as discrimination and exploitation are a part of the Justice and Fairness principle, which are basic human rights which might be violated through the misuse of data. There are different definitions of equality and fairness based on geographical location and culture.

Privacy and Data Protection															
Collection Limitation and Retention		Data Quality		Purpose Specification		Use Limitation		Transparency		Individual Participation and Access to Data		Anonymity		Individual Privacy	
A/D	Additional comments	A/D	Additional comments	A/D	Additional comments	A/D	Additional comments	A/D	Additional comments	A/D	Additional comments	A/D	Additional comments	A/D	Additional comments
Disagree	While it is true that companies are collecting personal data which is not really needed for their analysis, the collection and retention of data by itself is not unethical. The moment that some algorithm is applied on the data, based on the purpose of the analysis	Agree	If companies collect data from a user that is not aware of it, if some information is wrong in the data the user cannot give that company updated, correct information to them because they are not aware of that company having that information in the first place.	Agree	At the moment companies collect data it is not known yet what they will use the data for in the future: this makes it difficult for them to communicate it to the users. Companies should be sending some kind of warning that their data is being used for a certain purpose, which is not something they are currently doing.	Disagree	Companies are not allowed to share data with other companies; also, the company that collected the data is responsible for whatever is done to it (however the GDPR allows sharing as long as the intentions are clear and the sharing is done on a lawful basis). https://gdpr.eu/data-sharing-bounty-fine/	Agree	It is possible that even the developers that create algorithms do not completely understand how they work. Even if they did, it would be difficult for a company to explain their decisions to the public.	Agree	It is very difficult for users to participate in the process, since normally all the data is pushed towards the analytics.	Agree	Anonymity is quite difficult to enforce: it can be easy to trace back to a certain individual that matches certain characteristics from a data set. Also, combining datasets makes it easier to trace back to a certain person.	Agree	The matter of who has access to the data threatens the individual's privacy. Additional example: 1) Pregnant teenager whose father got to know she was pregnant because he received offers from the supermarket for baby products. 2) Somebody meeting a neighbor in the elevator congratulated them for their birthday because they were working at a bank and got to know through the bank's system when that person was born.

Agree	A company that has privacy as a principle will have to restrict itself from collecting too much data. Restricting the collection of data means potentially losing competitive edge, however privacy can be by itself the source of competitive advantage for a company.	Agree	Data quality might be affected by the intention of data scientists to get more interesting results - they might falsify records to get them.	Disagree	If customers are informed and they give their consent then a company can take their data, analyze it and see what results come out of it without facing ethical issues. There is no problem in having data and if there is the inform consent from the user it is fine to analyze it.	Disagree	There is no problem intrinsic to the sharing of data with other parties, but it is the action of enriching data after sharing that might expose the identity of a person and thus violate their informational privacy. If the information collected in a data set is enriched it is possible to be able to identify specific individuals from the data.	Disagree	The opaqueness of algorithms is only a problem if humans are extracted from the system, and thus they are no able to control the decisions anymore. The human is instead able to go back a decision made by the algorithm and make an ethical evaluation themselves.	Agree	At the moment it is still an unsolved issue because it is difficult to do.	Agree	/	Agree	The process of Big Data is often automated and insensitive, not only to privacy but to all kinds of ethical issues: a computer doesn't discriminate, it just says what the data is.
Disagree	/	Disagree		Disagree		Disagree		Disagree	Real transparency is not about posting more than 100 pages of policy notice on a website. There are very few companies that are truly transparent about their data.	Disagree	/	Disagree	/	Disagree	/

<p>Agree</p>	<p>Data gathering without limitation raises privacy concerns. It can be unethical to collect lots of data in the sense that there can be consequences for the users if some data is lost, or things happen to the data that the collector is not in control of (things they didn't plan for at the moment of collection).</p>	<p>Agree</p>	<p>If wrong conclusions are drawn from wrong data, this might impact something unexpected or unwanted, which can ultimately impact people.</p>	<p>Agree</p>	<p>Not having an hypothesis for what the data will be used for, or using the data for a different purpose than the one initially intended at the moment of the collection impacts the user's privacy negatively and is unethical. Furthermore, users are often not informed correctly and most people do not understand what happens to their data.</p>	<p>Agree</p>	<p>Whenever data is being shared between companies, the user loses control of it, which makes it difficult to impose the Right to be Forgotten: if the user is unaware of which company holds their data, they cannot ask them to delete it. Every company with data should send out the consent agreement to the third party prior to sharing data.</p>	<p>Agree</p>	<p>When the algorithm is a black box it can raise privacy, and therefore ethical concerns. A black box algorithm cannot be explained, and especially in the context of governments (as well as other companies) transparency is fundamental because their decisions have to be explainable.</p>	<p>Agree</p>	<p>Users have the right to participation but it is very difficult to enforce. There should be a third party that checks and audits each important application of Big Data to make sure that the user data is being used according to the terms of the consent agreement : this way users can be informed of whether their data is being treated correctly.</p>	<p>Agree</p>	<p>When a user is promised anonymity of their data but there is enough information in the data to make the identification possible, a wrong promise was made and such behavior is unethical. However, the possibility of de-anonymizing a person is less unethical than intentionally trying to identify a person from a data set.</p>	<p>Disagree</p>	<p>There are measures to prevent privacy issues in Big Data, which would make it in turn less sensitive to these issues. It is not directly unethical to have Big Data systems in place, and there are countermeasures for these problems.</p>
--------------	---	--------------	--	--------------	---	--------------	--	--------------	---	--------------	--	--------------	--	-----------------	--

Agree	This principle is also demanded by the GDPR, which tries to prevent unnecessary data to be collected. This is in sub-5, which says that companies can only ask to collect data on the data subject which is relevant for its purpose.	Agree	/	Agree	Not knowing the results of the analysis makes it difficult to know the purpose of collection in advance.	Disagree	The purpose limitation means both purpose specification and use limitation. The use limitation concerns the processing of data for a particular purpose.	Agree	Transparency is mentioned in the GDPR: "data must be processed lawfully, fairly and be transparent". In Article 12 the details are further elaborated: "the controller must take measures to provide information from the following articles to the data subject". <i>Additional example:</i> Syri is a governmental system used to detect social security fraud, which however often uses data from a poor neighborhood and combine it with other data to point out which people are more likely to commit fraud.	Agree	Organizations are obliged to comply with the users' requests. The GDPR, however, does not give users the right to participate in the process, but only control over data (such as the right to correct incorrect data or to delete it)	Agree	The GDPR imposes a very high threshold for data to be anonymous: removing a name and an address is not sufficient. Investigations have shown that even when identification chances are deemed to be very low, it is not quite the case when dealing with Big Data.	Disagree	Individual Privacy is hard to understand as a principle because it already belongs to the Privacy category, which already discusses the topic in broad terms.
Disagree	/	Disagree	/	Disagree	/	Disagree	/	Disagree	/	Disagree	/	Disagree	/	Disagree	/
Disagree	/	Disagree	/	Disagree	/	Disagree	/	Disagree	/	Disagree	/	Disagree	/	Disagree	/

Agree	/	Agree	The labelling of data is done by somebody else, who might be affecting its quality. On one hand there is the problem of profiling taking place on the basis of bad quality data; on the other hand it is hard to determine whether the collection of the data was in the first place just.	Agree	The problem exists, however the GDPR is very specific about how an organization can contract and subcontract the use of data: it should be illegal to use data for a purpose different than the one specified but it probably happens all the time. The goal of the collection should not change throughout the data lifecycle.	Agree	The unethical share of data occurs even within the same organization. <i>Additional example:</i> with anti-laundering efforts marketing and compliance work together when clients onboard. They ask questions to clients and collect personal data about them, however it is ambiguous how that information will be used and the data subject is probably unaware of it.	Agree	Algorithms are progressively going to become more complex and more self-thought, so we will probably never reach a point in which we will not really understand it. A countermeasure for the opaqueness of algorithms is the integration of the human in the process, in a way that the human can do arbitrage and challenge the algorithm, or even do a do-over; this would also allow the integration of ethical principles in the data process.	Agree	The GDPR gives power to the users to reclaim their data from companies, but there haven't been any interesting examples where they have.	Agree	/	Disagree	Individual Privacy overlaps with other categories.
-------	---	-------	--	-------	---	-------	--	-------	--	-------	--	-------	---	----------	--

Annex J

		Big Data Activities	Data Governance Practices
Lifecycle Phases	Data Collection	Data collectors not asking users for consent in a specific, informed and unambiguous way.	Accountability for the data collection. Openness towards customers and authorities.
		Excluding individuals from the data, causing the data sets to not be representative of the population and individuals not receiving the right offers.	Reducing error of data by looking for biases in the way data is collected.
		Data gathering without limitation having consequences on the user (e.g. if the data is lost)	Defining a data strategy to take control of the data collection.
		Sharing data within and outside a company, causing the user to lose control of it.	Accompanying data with information about its provenance and processing. (Kroll, 2018) Evaluate data for fidelity to the phenomenon under consideration. (Kroll, 2018)
		Collecting personal data can affect customer trust and the company's reputation, and puts the company at risk of legal noncompliance. (Kroll, 2018)	Designating a review board responsible for approving or denying the collection of new data. (Kroll, 2018) Establishing effective policies and procedures to guarantee alignment between business principles and the collection of data.
		Use of algorithms that manipulate the user's decisions and limits their autonomy.	Introducing standardized operations to evaluate algorithms.
	Data Analysis	Predictions inferred from flawed Big Data models putting the safety of users in danger.	Establishing a common understanding of algorithms.
		Use of algorithms to target individuals in a personalized way.	Reducing error of analysis.
		Big Data algorithms that take discriminating characteristics into consideration putting users into a position of disadvantage.	Looking for systematic biases in the way outcomes are labelled, outliers are pruned, groupings are defined and categorical variables are encoded. (Kroll, 2018)
		Humans designing algorithms with their own perspective in mind introducing errors and biases that can harm users.	Accountability of data usage.
		Using black box algorithms that are difficult to understand and explain to the user.	Establishing effective policies and procedures to guarantee alignment between business principles and data analysis. Defining a data strategy to clarify how the company works with data.
		Investigating sensitive questions using company data. (Kroll, 2018)	Designating a cross-functional review board responsible for examining the details of data analysis. (Kroll, 2018)

		Sending individuals personalized offers and information, creating a situation of unfairness.	Monitoring the performance of a system after launch by means of black box testing to test it against unfairness. (Kroll, 2018)
		Drawing conclusion from poor quality data, with repercussions on the data subjects.	Designating an internal role responsible for owning the outcomes of analysis. (Kroll, 2018)
		Users being unaware of how predictive information is inferred and impacts them.	Establishing a common understanding of how specific decisions are made (transparency).
		Generating mistaken decisions and false positives, causing the putatively high-risk individuals to be treated unfairly. (Kroll, 2018)	Considering the possibility that mistaken decisions might disproportionately harm individuals or protected groups and testing for feedback loops. (Kroll, 2018)
		Using data-driven systems, which can be affected by modeling errors and whose fidelity changes over time. (Kroll, 2018)	Having data scientists validating predictions and monitoring the performance of systems after launch. (Kroll, 2018)
		Deployment of insights from analysis of sensitive questions. (Kroll, 2018)	Establish effective policies and procedures to guarantee alignment between business principles and decision making phase Designating a review board responsible for approving or denying the use of analytics insights. (Kroll, 2018)
Data Cleansing		Data scientists affecting data quality with the intention to get more interesting results.	Looking for systematic biases in the way data is cleansed and validate cleansing assumptions. (Kroll, 2018)
		Data scientists choosing how to describe data and missing details in the world. (Kroll, 2018)	Validating assumptions baked into the normalization methodology. (Kroll, 2018)
			Establish effective policies and procedures to guarantee alignment between business principles and the data cleansing phase. Define a data strategy to clarify how the company cleanses data.
Data Visualisation		Humans interpreting analysis results introducing errors and interpreting the results based on their personal values.	Being aware of human bias which affects the way results are interpreted.
Data Archival and Data Deletion		Reidentification of user data despite promising users anonymity of their personal information.	Considering the risk that retained data could be reidentified, which depends on the type of data in question and the context in which it is being used. (Kroll, 2018)
			Understanding how and why data must be retained and how it will be used, to know how it can be profitably discarded or properly anonymized to minimize risk. (Kroll, 2018)

Annex K

Concept 1	Concept 2	Additional comments
Autonomy	Dignity	Autonomy and Dignity both refer to the same example.
		Dignity is the underlying principle of autonomy.
Autonomy	Individual Privacy	It is helpful to have privacy rules to make sure that informed consent leads to the autonomy of users.
Autonomy	Informed Consent	Informed consent has a relation to autonomy, specifically to what is referred to in the autonomy definition as 'meaningful choice'.
Dignity	Discrimination	If the data set is misused to discriminate on factors such as gender, age, etc., then a violation of dignity is happening.
		When there is a legal case of discrimination, dignity is violated; however not in all cases where dignity is violated it is necessarily a case of discrimination.
Dignity	Equality and fairness	Dignity is the underlying principle of fair treatment and justice.
Safety	Discrimination	Biases in a Big Data system might cause it to put the safety of people at risk.
		Profiling people in the wrong way can put their safety at risk.
Safety	Data quality	Incorrect data can affect the safety of individuals.
Data quality	Social solidarity, inclusion and exclusion	If data is wrong a user might be excluded from getting offered certain services.
Data quality	Individual Participation and Access to Data	The difference in knowledge of the population can cause representativeness problems in the data sets: higher educated people will be more aware of what is done with user data and are more likely to decide to not participate in the process.
Imbalance of power	Informed Consent	Informed consent sounds more like a situation of imbalance of power in which the data collector leaves the user in the dark, does whatever they want with the user's data and may inform them with privacy notices but the user might not have a clue about what the notice is really saying
Imbalance of power	Autonomy	Power is autonomy, and the power of a person to be in their own control.

		Data subjects controlling the usage and outcome of data can create more balance in the power relation with companies.
Imbalance of power	Transparency	The lack of transparency in the way data is collected and used creates an imbalance of power.
		Transparency is seen as something that can correct the imbalance of power, because it gives people the possibility to react.
Imbalance of power	Individual Participation and Access to Data	/
Equality and Fairness	Discrimination	/
Equality and Fairness	Value sensitive design	The algorithm is never intrinsically unfair, but that is always because of the human who made it.
Collection Limitation and Retention	Purpose Specification	Not knowing the relationship between the goal and the data, in other words the purpose of the collection, it is difficult to minimize the data collection, as it might be only obvious afterwards that certain data was unnecessary.
Collection Limitation and Retention	Informed Consent	Informed Consent is what organizations use to justify the collection of data beyond the legal limits.