



**POLITECNICO  
DI TORINO**

Master degree course in Engineering and Management

Master Degree Thesis

# **The economics of Machine Learning: a microeconomic model of customer-firm interaction**

**Supervisor**

prof. Carlo CAMBINI

**Candidate**

Matteo Pio BORTONE

student ID: 241956

December 2018



# Summary

Alan Turing, who had a pivotal role in today's computer science, stated in the 50s: "I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted."

A little less than twenty years since the beginning of the twenty-first century, the Artificial Intelligence (AI) field is in ferment and continuous expansion: Turing's vision lives in the work of researchers, academics, professionals, who try to push the limits of the current knowledge and to contribute towards the creation of an AI comparable to a human intelligence.

Although this is a matter related more to Information Technology than to other sciences, AI is a technology with virtually unlimited potential is also fascinating and attracting the attention of scholars from the economics and policymaking sphere.

Will AI take over the hard and dangerous work? Will economy prosper because of AI? Will consumers be deprived of their privacy? Will a few, enormous monopolists thrive at the expenses of the mass? To make sure that AI bolsters collective progress, this and many other questions need to be addressed. This thesis is a small, shy step in this direction.

Because of their tight bond with data, AI and ML will result in a change in customers-firm interplay, making necessary to address the economic impact of these technologies, so that policymakers can act to protect the customer's interest, especially concerning data sharing. The economy of Machine Learning reduces to the economics of data, in which smart algorithms are just a fresh way to look at how to extract value from the information that consumers share about themselves.

This thesis is an attempt to create a microeconomic model of customer-firm interaction, in the case of a firm offering a product featuring Machine Learning properties and the customer sharing data and personal information, an essential ingredient in the creation of value from smart algorithms, with the firm. The customer-firm interaction is therefore modeled to take place in two interconnected markets: one for the product and the other for the customer data. The firm is assumed to be a monopolist first and a social optimizer later, to highlight the differences for consumers welfare.

The remainder of the document is structured as follows.

The first chapter introduces some of the most debated international topics on AI and how it could shape the society of the future, from both a technological and an economic perspective. A focus on the macroeconomic impact on productivity and employment follows, and the chapter is concluded with a digression about the reasons why the AI is a concern for policy-makers, for either economics or other fields.

The second chapter is about to the microeconomic model of customer-firm interaction, organized in order to introduce the assumptions under which the model holds first, and explore the

behavior of the consumers and the firm (first a monopolist, then a social optimizer) in the defined environment then. The chapter ends with a comparison that clarifies which solution (under which conditions) is to be preferred from a consumer perspective.

The third chapter consists of two sections. The first part presents two case studies related to two companies that use ML techniques in their products, to understand to what extent the presented model is applicable in real situations. The second part is a normative exposition built upon the model presented in the second chapter, aiming at providing cues for the policymakers interested in protecting the customers from a possible power unbalance with the firm. The topics discussed are the effects of bias on consumers, the problem of consumer manipulation and data privacy.

The fourth and last chapter is the conclusive one that summarizes the main findings of the research work and gives suggestions for future research paths, either being a natural evolution of the model or stemming from it.

The appendix reported at the end of the document contains the mathematical proofs of the results presented in the second chapter of this thesis.

# Acknowledgements

In preparation of my thesis, I had to take the help and guidance of some respected persons, whom all deserve my deepest and most sincere gratitude.

Nobody has been more important to me in the pursuit of this project than the members of my family. I would like to thank my parents and my brother, whose love and guidance are with me in whatever I pursue. They are the ultimate role models.

To my friends and roommates, thank you for listening, offering me advice, and supporting me through this entire process. May the connection we have never be lost.

# Contents

<b>List of Tables</b>	IX
<b>List of Figures</b>	X
<b>1 Introduction to Artificial Intelligence</b>	1
1.1 Technological insights on Artificial Intelligence . . . . .	2
1.1.1 Enabling factors for AI . . . . .	2
1.1.2 Definition of AI . . . . .	3
1.1.3 The building blocks of AI . . . . .	4
1.1.4 Current state of the art . . . . .	6
1.1.5 AI and games: a proficient match . . . . .	7
1.1.6 Current limits and future developments . . . . .	8
1.2 The pervasiveness of AI, domain by domain . . . . .	10
1.3 The relevance of AI in Economics . . . . .	11
1.4 Macroeconomic impact of Artificial Intelligence . . . . .	12
1.4.1 A regional breakdown . . . . .	13
1.4.2 Impact of AI on productivity . . . . .	16
1.4.3 Impact of AI on employment . . . . .	18
1.5 Policy implications of AI . . . . .	23
1.6 Microeconomic dynamics brought by AI . . . . .	24
<b>2 A microeconomic model of customer-firm interaction</b>	27
2.1 ML: an economics perspective . . . . .	27
2.2 Model scope . . . . .	28
2.3 Model structure . . . . .	29
2.4 Two-sided market setup . . . . .	29
2.5 Consumer's behavior . . . . .	30
2.5.1 Consumer-related variables . . . . .	32
2.5.2 Discussion of model assumptions . . . . .	33
2.6 Firm's behavior . . . . .	34
2.6.1 Firm-related variables . . . . .	34

2.6.2	Number of sub-games	36
2.7	Monopolist firm's behavior	36
2.7.1	Generic stage game of the infinite sequence: optimal values for $p, d$	37
2.7.2	Stage 1: optimal value for the investment	38
2.7.3	Consumer surplus	38
2.8	Social optimizer's behavior	39
2.8.1	Generic stage game of the infinite sequence: optimal values for $p, d$	39
2.8.2	Consumer surplus	40
2.8.3	Stage 1: optimal value for the investment	41
2.8.4	Generalization of the total Welfare function	41
2.9	Solution comparison: Monopolist VS Social Optimizer	41
2.10	A numerical example	42
2.10.1	Monopolist solution	43
2.10.2	Social Optimizer solution	44
2.10.3	Solutions comparison and discussion	45
2.11	Chapter conclusions	46
<b>3</b>	<b>Case study discussions and normative recommendations for policymakers</b>	<b>47</b>
3.1	Case study of two digital platforms	47
3.1.1	Digital video streaming: Netflix Inc.	48
3.1.2	Online retailing: Amazon Inc.	49
3.1.3	Cases discussion	49
3.2	Issues arising from smart products and relevant for the regulator	51
3.3	Issues related to bias in smart products	51
3.3.1	Nature of the problem	52
3.3.2	Recommendations for the policymaker	54
3.4	Impact of AI recommendations over consumers' decisional power	55
3.4.1	The choice problem and choice overload	56
3.4.2	The problem of manipulation	57
3.4.3	Recommendations for the policymaker	58
3.5	Data Privacy	58
3.5.1	Taxonomy of data	59
3.5.2	Problem of identifiability	60
3.5.3	Data as a consumer policy law matter	61
3.5.4	Recommendations for the policymaker	61
<b>4</b>	<b>Conclusions and further developments</b>	<b>63</b>
4.1	Future developments of the model	63
4.2	Future research based on the model	64

<b>Appendix - Mathematical proofs</b>	75
Derivation of the demand functions - equations 2.3 . . . . .	75
Derivation of the number of the sub-games - section 2.6 . . . . .	75
Derivation of the monopoly solution - equations 2.13 . . . . .	76
Derivation of the consumer surplus in case of a monopolist firm - equation 2.18 . . . . .	77
Derivation of the welfare maximization solution - equations 2.23 . . . . .	77
Derivation of the welfare maximization solution - equations 2.29 . . . . .	79
Derivation of the consumer surplus in case of a social optimizer firm - equation 2.18 . . . . .	80
Derivation of the preference conditions of social optimizer versus monopolist solution . . . . .	81
Price comparison - equation 2.30 . . . . .	81
Quantity comparison - equation 2.31 . . . . .	81
Shared data - equation 2.32 . . . . .	82



# List of Tables

1.1	GDP changes by country by 2030. Depiction of The Economist Intelligence Unit (2017) . . . . .	13
1.2	Average annual TFP growth in selected countries, 1985-2017. Source: <i>OECD</i> . . .	19
1.3	Market capitalization VS number of employees; selected companies. Source: <i>ycharts.com</i> , <i>forbes.com</i> . . . . .	20
2.1	Monopoly VS Social Optimizer solutions . . . . .	42

# List of Figures

1.1	Annually Published AI Papers, by year (Shoham et al., 2017)	3
1.2	Active startups developing AI systems (USA), by year (Shoham et al., 2017)	3
1.3	Annual VC investment in AI startups (USA), by year (Shoham et al., 2017)	4
1.4	Number of transistors per microprocessor, 1970-2017; depiction of Rupp (2018)	4
1.5	Millions of Floating Point Operations per Second, 1993-2017; depiction of TOP500 Supercomputer Database (2018)	5
1.6	Microprocessor clock speed (Hertz), 1975-2016; depiction of Kurzweil (2018)	5
1.7	Cost of memory (\$/Mbyte), 1957-2018; depiction of McCallum (2018)	6
1.8	Accuracy of the best AI in object detection, by year (Shoham et al., 2017)	7
1.9	Accuracy of the best AI in speech recognition, by year (Shoham et al., 2017)	7
1.10	Accuracy of the best AI in question answering, by year (Shoham et al., 2017)	8
1.11	Innovation environment (on the left) and technological readiness (on the right) in large advanced economies and large emerging economies, 2009-2017 (Schwab and Sala-i-Martin, 2017)	14
1.12	World GDP from 1700 to 2016; The vertical lines denote the starting point of each industrial revolution. Source: <i>The World Bank</i>	17
1.13	Real GDP (constant 2010 US\$) growth rate, by five-year periods and by Country. Source: <i>The World Bank</i>	17
1.14	Growth of Capital Services provided by ICT assets, selected Countries, 1990-2016 (The Conference Board, 2017)	19
1.15	Evolution of labor market flexibility by region (on the left) and within the European Union (on the right), 2007-2017 (Schwab and Sala-i-Martin, 2017)	20
1.16	Percentage of jobs at risk of automation, by 2030, by geographical region and industry (Gillham et al., 2018)	21
1.17	Share of jobs requiring AI skills on the portal Indeed.com (selected Countries), by year (Shoham et al., 2017)	22
1.18	Job Openings, disaggregated by required skills on the portal Monster.com, by year (Shoham et al., 2017)	22
1.19	Classification of jobs according to the Axis of creativity, as proposed by Kai-fu Lee.	23
2.1	ML from a technological and from an economics perspective.	28
2.2	From data to prediction: logic process depiction.	28
2.3	Model breakdown structure.	29
2.4	Depiction of the consumer-firm interaction, including the direct and cross-network effects.	30

2.5	Trend of the variables $U_{ML}$ and $c_{BIAS}$ versus the stock of collected data . . . . .	34
2.6	Trend of the Machine Learning efficiency versus the stock of collected data . . . . .	36
3.1	Example of biased AI: a dog (panel b) misidentified for a wolf (panel a) because of the similarity of the background . . . . .	52
3.2	Depiction of a biased route (red line) suggested by a web mapping service instead of the optimal one (green line). . . . .	53
3.3	Example of data classification according to their sensitivity. . . . .	60

# Chapter 1

## Introduction to Artificial Intelligence

The human history is full of discoveries, inventions, and revolutions that, by improving automation and connectivity, led to wealthier and more productive societies. Starting from the invention of the steam engine in the 18<sup>th</sup> century, the electricity in the 19<sup>th</sup> century and the advent of electronics in the 20<sup>th</sup> century, people is currently encountering a fourth industrial revolution that gravitates around Artificial Intelligence (AI): if the former three were about generating artificial power, the latter is about artificial smartness. This outline requires, to anyone involved in promoting and developing the AI field, an extended consideration on the real capabilities of this technology to deliver new sources of long-term prosperity, as the very way of being human could change drastically.

More generally, expectations on the consequences of the fourth industrial revolution are high, and this is expected to strike at least three distinct areas (Baweja et al., 2016): (1) the technosphere, since AI involves manifold scientific disciplines and engineering; (2) the natural world, since people are now capable of monitor, examine and digitize anything related to natural phenomena at scale and a very high speed and precision; (3) finally, the fourth industrial revolution will impact the human world because agents endowed with human-like intelligence enable new ways of connecting, interacting with other people and process information.

Among all, this last point is particularly important: live in a very complex world as today's could cause people to drop the effort to seek information just because there is too much to acknowledge: this, in turn, leads to taking less educated decision; from an economic point of view, this is everything but a good thing. A fundamental principle of economics is the optimization, that is making the best choice according to the information that an agent has at a given time; reduced, or absent information undermines the optimization process, resulting in a lower gain (if not a loss) for the agent. For the doer to reach a state of equilibrium, it is necessary to process all the information available, possibly using tools like AI. This choice grants its users an advantage of volume (the world we live in is too complicated for us to handle all the knowledge just with human brains) and advantage of time (a person can gather information at lightning speed, often also presented in natural language).

This chapter presents some of the most debated international topics on AI and how it could shape the society of the future. The starting point is a technological focus on what AI is; after having equipped the reader with fundamental knowledge about the potential offered by smart machines, the discussion moves on to a purely economic discussion. After having motivated the relevance of AI for the economy, the study proceeds with a review of the macroeconomic expectations of the AI (for regional base first, concerning productivity and employment then). The chapter closes with a preface to the influence of these technologies on the consumer and, consequently because these also demand the proper consideration of the policymaker.

## 1.1 Technological insights on Artificial Intelligence

### 1.1.1 Enabling factors for AI

AI is a technology that stems directly from Information Technology and, hence, shares similar properties of pervasiveness; to date, it can be observed in many products, services, and decision-making processes. Its rise in the past decade is due to three enabling factors: data abundance, a determined commitment of researchers and entrepreneurs, the technological advancement of computers. All of them contributed to driving down the cost of these solutions and above all, made them more accessible to design, program and implement.

First, the abundance of big data acted as the rocket fuel to the AI engine. The AI - big data relationship is bi-directional: big data relies on AI to extract information while, at the same time, algorithms need data to be trained and to perform their tasks with proper accuracy. If intelligent machines cannot rely on a sizable amount of data to draw information from, they cannot improve themselves or be smart at all. AI applications are only as robust as the depth and quality of the data behind them, and this is why these two technologies have advanced in lock-step in the past few years. To put things in perspective, it has been estimated that people produce 2,5 Exabyte<sup>1</sup> of data everyday (Marr, 2018), the 90 percent of which has been generated in the past two years alone. Although the amount of data needed for machine learning purposes depends both on the complexity of the problem and on the complexity of the chosen algorithm, the sizes involved are usually massive. For instance, it takes more than 4,000 pictures of Computed Tomography to train a highly accurate neural-network based classifier (Cho et al., 2015) and more than 200,000 videos to train another neural network to classify sports videos according to their content (Karpathy et al., 2014). Apart from abundance, data quality, speed and variety are equally significant. As a matter of fact, nowadays data are no longer collected in the form of symbolic data, but they are acquired in the form of images, footages, audio, and many other forms so that they can better represent diverse phenomena; algorithms can, therefore, rely on their higher level of details and offer better insights.

Secondly, fresh approaches to the puzzle of coding intelligence made it easier to program these systems. Even though the field of AI was already founded as an academic discipline back in 1956, it is only in the past decade that it made astounding advancements. According to the Scopus database of academic papers<sup>2</sup>, the number of Computer Science publications about AI has increased by more than nine times since 1996, going from about a thousand to over 19 thousand in 2015 (Shoham et al., 2017). A comparable growing trend can be found in the number of active venture-backed US private companies developing AI systems, that moved from almost zero in the first five years of the 90s to over six hundred in 2016; the number started growing exponentially since 2010. This trend ultimately reflects the growing annual venture capitalists' investment in AI startups, which has expanded six times since 2000 (Shoham et al., 2017).

The third (and last) enabling factor for AI is associated with the technological progress in electronics and computer science, that led to more computational power at a lower price. As depicted in figures 1.4, 1.5, 1.6 and 1.7, the computational power of computers rose exponentially, reaching performances that are thousands of times higher than a few decades ago. Such progress, along with a terrific contraction in the cost of memory and more power efficient machines made it possible the design of powerful supercomputers capable of process tremendous amounts of real-time data and work in connection with other supercomputers. The current state of technology, albeit very advanced, is still lacking the proper power required to build smarter AI. Cloud computing can only be an answer in the short time: what is needed are essentially more powerful stand-alone machines, eventually based on a new knowledge base like quantum computing.

---

<sup>1</sup>Equivalent to 1 billion Gigabytes or  $10^{18}$  bytes

<sup>2</sup>The Scopus database contains over 200,000 papers in the field of "Computer Science" that have been indexed with the key term "Artificial Intelligence" and almost 5 million papers in the subject area "Computer Science"




Figure 1.1. Annually Published AI Papers, by year (Shoham et al., 2017)




Figure 1.2. Active startups developing AI systems (USA), by year (Shoham et al., 2017)

### 1.1.2 Definition of AI

Before proceeding further in the themes investigated in this thesis, it is useful to provide a few definitions to fully-understand what AI is and, on a very high level, how it operates.

There is no univocal definition of Artificial Intelligence because intelligence is itself a poorly defined concept and there is no general agreement about what can be defined as such. It can be referred to as “the theory and development of computer systems able to perform tasks normally requiring human intelligence” (Oxford Dictionary, nd), or as an amalgam of science and computational technologies that take inspiration by how humans use their nervous systems to percept the environment and take consistent actions (Stone et al., 2016). AI is nonetheless very different from human intelligence since the first is extremely specialized to work in a narrow domain, while the second is capable of sensing various inputs from a broad spectrum of sources and elaborate more flexible responses.

Though, if intelligence is defined the quality that enables an entity to function appropriately and with foresight in its environment (Nilsson, 2009), it follows that human intelligence is the real benchmark for AI. General Purpose AI is a big deal, consisting of a complex system that exhibits intelligent behavior across several domains, making it capable of performing very different tasks. Create such intelligent devices require an in-depth cross-disciplinary knowledge in computer science, mathematics, engineering, linguistics, philosophy, neuroscience; it is, however, improbable that machines will show signs of broadly-applicable intelligence in the next 20 years (The White House, 2016): the more context a task requires, the less likely a computer will be able to do it soon.

One of the most remarkable breakthroughs in the AI space is Machine Learning, a subfield of computer science that strives to build algorithms capable of learning from and make predictions on data (Samuel, 1959). By using Machine Learning, programmers can write simpler programs that do not require to specify how to react to every single input: put in simpler terms, with Machine Learning computers program themselves, by starting with a training set of data used to derive rules or procedures applied consequently (The White House, 2016). Advanced ML techniques




Figure 1.3. Annual VC investment in AI startups (USA), by year (Shoham et al., 2017)




Figure 1.4. Number of transistors per microprocessor, 1970-2017; depiction of Rupp (2018)

like Deep Learning and Neural Networks use complex networks of computers that resemble the human brain and its structure made of neurons organized on different layers. These technologies can convey a considerable amount of data within one single domain and learn to predict or decide at superhuman accuracy.

Another challenging application of AI aims to design systems capable of learning by interacting with both humans and data (whether structured or unstructured): this is what practitioners define cognitive computing (Dalton et al., 2015). The real power of cognitive computing relies on providing responses that are contextualized, quick and with a very high level of confidence.

Given the purposes of this thesis, which are more economics-focused than technology-focused, the terms AI, AI system, intelligent machine, smart system or smart algorithm are employed interchangeably to indicate a technology that manifests (to some extent) some form of intelligence and that is consequently capable of learning from data.

### 1.1.3 The building blocks of AI

The human brain is a highly complex organ, made up by many parts different in size and structure that perform different functions. Similarly, those involved in creating complex systems that mimic the functionality of a brain must decompose the problem into many subproblems (or components) specifically designed to perform a particular function.

This approach makes it is reasonable to think about a generic Artificial Intelligence as a combination of ten building blocks (or a subset of them) (Gerbert et al., 2017). The first three blocks are those capable of sensing the environment and extract data from it; the second three building blocks then interpret these data in order to let the last four ones to act consequently.




Figure 1.5. Millions of Floating Point Operations per Second, 1993-2017; depiction of TOP500 Supercomputer Database (2018)




Figure 1.6. Microprocessor clock speed (Hertz), 1975-2016; depiction of Kurzweil (2018)

Not surprisingly, many researchers working on AI focus on just a few of these parts. The list below provides a brief description for each of the building blocks mentioned above.

**Machine vision** This block classifies real-world objects starting from images, video recordings or other signals.

**Speech recognition** It is implemented to obtain text starting from spoken words.

**Natural-language processor** It helps to detect the intent of the interlocutor in text-based commands.

**Information processing** It is a comprehensive block that includes methods of search or knowledge extraction to provide answers to queries.

**Learning block** This block is what has previously been defined as Machine Learning.

**Planning and exploring agents** It helps the AI to identify the best sequence of actions to achieve a goal.

**Image generation** It gives the AI the ability to generate pictures based on models.

**Speech generation** It gives the AI speech capability in order to communicate with humans.






Figure 1.7. Cost of memory (\$/Mbyte), 1957-2018; depiction of McCallum (2018)

**Handling and control** It makes the AI capable of handling and interacting with real-world objects.

**Navigating and movement** This block allows the AI to move safely through its environment, avoiding objects and obstacles.

For instance, the well-known Virtual Personal Assistants (VPA) Siri (by Apple) or Cortana (by Microsoft) use speech recognition and natural-language processing algorithms to convert the user's input in data that the software can understand; then the information is processed on Siri Servers, where a proper response is generated and conveyed back to the user through speech and text generation. The transaction is also handled by machine learning algorithms so that the VPA can learn from it.

#### 1.1.4 Current state of the art

The intent of building thinking machines is without any doubt, an ambitious and audacious endeavor; start-ups, corporations, and universities are forges of ideas and solutions that, even when tiny and apparently not relevant, could result in a cumbersome step for the gain of the future society. The AI field is so abuzz that taking a snapshot of the contemporary state of the art can represent an already-old situation after a short time. However, to appreciate how effectively smart machines can influence the daily life of a man, the state of technical performances in specific domains is depicted nevertheless. They all represent major technological milestones and find various employment in several fields.

Computer vision is an essential feature of many robots. Even autonomously driven cars, despite the ability to implement highly sophisticated GPS tracking systems, would not make sense if unable to perceive other vehicles or obstructions on their way. Depending on the scope of the AI, vision can be broken down in several general tasks, such as object recognition, motion analysis, and position-and-orientation estimation. Fueled by the latest advances in Deep Learning, the best AI system has already overtaken human performance since 2014, in a Large Scale Visual Recognition Challenge Competition (LSVRC). Recently, object recognition has become quite mainstream, since it can easily be found in many web applications. Google's Lens app allows its users to explore what they have around by merely pointing their camera: it makes possible to learn about famous landmarks, to translate written text, to identify plants and animals. Another commercial application of computer vision is face recognition, which is mainly deployed for security features on IT products and services. Despite this, computer vision has still hard times in dealing with poor quality pictures, adverse lighting, low resolution, and tricky camera angles.

A different AI reached, in 2016, human-level accuracy in speech recognition from phone call audio, becoming over 10% more accurate in just five years; the most advanced solutions are also capable of recognizing many languages and related variants (eventually by automatically recognizing which one is being spoken) and filtering inappropriate language. Performances about question answering made a huge leap forward in just two years, passing from the accuracy of 60% to almost 80%, slightly below human accuracy (Shoham et al., 2017). In real cases, however, the accuracy of the AI depends on the openness of the domain of the possible questions (AI generally perform better when the domain is narrower), since a correct performance requires a deep semantic understanding, meaning that the AI needs to be able to perform complex anaphora resolution, or use common sense as people usually do. Once again, it is quite challenging to translate these human-like features in algorithms.




Figure 1.8. Accuracy of the best AI in object detection, by year (Shoham et al., 2017)




Figure 1.9. Accuracy of the best AI in speech recognition, by year (Shoham et al., 2017)

### 1.1.5 AI and games: a proficient match

Since the ideation of AI as a science, games have been an excellent way to assess the capabilities of AI. They are especially convenient for testing the capacity of an AI because they make accessible to quantify performances through numeric scores and win-lose outcomes.

In 2016, AlphaGo, an AI created by DeepMind (currently, an Alphabet/Google subsidiary) defeated with a score of 4 matches to 1 Lee Se-dol, the world champion of the old Chinese game Go (Devlin, 2018). This achievement is quite astounding since Go has  $250^{150}$  combinations, and it is therefore impossible to be played and won by using a brute-force approach. The system used a hybrid of AI techniques: its creators partly programmed it, but it also taught itself using deep reinforcement learning. One year after the achievement, DeepMind published a paper in which the authors explain how they created a new version of AlphaGo, AlphaGo Zero, that works without possessing any previous knowledge of the game, and that he, therefore, managed to master the game by being the master of himself (Silver et al., 2017).

Board games, however, present limitations. First, they are turn-based, which means that the AI is not forced to make decisions in a constantly changing environment. Second, the AI has access to all the information in the environment and doesn't have to make guesses or take risks




Figure 1.10. Accuracy of the best AI in question answering, by year (Shoham et al., 2017)

based on unknown factors. These reasons justify why the attention of the researchers shifted towards more complex games, including video games.

In April 2017, an AI developed by Carnegie Mellon University researchers defeated four human players in a Texas Hold'em competition held in Hainan, China (Spice, 2017). Mastering even a two-player form of poker is an astounding achievement for AI because poker requires players to act with limited information, and to sow uncertainty by behaving unpredictably. To win, both instinctive judgment and caution are necessary, but these qualities that do not belong to a computer. Lengpudashi (this is the name given to the CMU AI) won by using a game-theory-based algorithm, which could be very useful in many other applications, like financial trading and business negotiations (Knight, 2017).

To date, many AIs can master many games developed in the '70s and '80s, and some of the most advanced are approaching games from the 90s. Researchers from OpenAI found a way to make the AI do something, without expressly telling it what its goal would be: the curiosity of the artificial agent purely drove the entire process of discovery. By playing Atari Games, Super Mario Bros., experimenting with virtual 3D navigation and even the Pong game played between two competing AIs. The curiosity-driven agent kind of sets its own rules motivated to experience new things. When it plays Breakout - the brick-breaking game - it performs well because it does not want to get bored: "The more times the bricks are struck in a row by the ball, the more complicated the pattern of bricks remaining becomes, making the agent more curious to explore further, hence, collecting points as a by-product. Further, when the agent runs out of lives, the bricks are reset to a uniform structure again that has been seen by the agent many times before and is hence very predictable, so the agent tries to stay alive to be curious by avoiding reset by death" (Burda et al., 2018).

### 1.1.6 Current limits and future developments

Many fields have, in the past, ran into their fundamental limits, both practical and theoretical (in physics, it is not possible to accelerate faster than the speed of light). To consider if there is anything to prevent humanity to reach the desired endpoint in AI development, it is first of all necessary to define it.

For smart machines, the ultimate ending point is represented by strong AI, that refers to such a powerful AI that possesses consciousness, self-awareness, emotions, and morality; despite the importance that such achievement would have, this would come with many ethical problems, such as whether strong AI should have rights, or whether they should be allowed to switch themselves off. Machines that think are indeed a provocative (and at times disturbing) idea, so many researches and famous entrepreneurs oppose even to the idea of building a strong AI. Instead, they usually support a less extreme endpoint, but certainly more easily accessible, represented by Artificial General Intelligence, a machine that comes with all the benefits of a strong AI but without consciousness, emotions and all the ethical problems related.

In the end, AI does not need to be perfect. They need nothing but be better than humans; in an ever-increasing number of cases, this is already happening. That said, AI is still immature

under several aspects, and the path towards a General AI is not free from obstacles: tasks for AI systems are usually framed in restricted contexts for the sake of making progress on that specific area. While machines may exhibit astounding performances on a specific task, performance may degrade dramatically if the task is altered even slightly. In order to meet the expectations that society has on AI, it is important to focus the research efforts and devote resources to (Russell et al., 2015):

- Enhance AI robustness, especially its verification (ensure that an evolving system like AI keeps working in a proper way), validity (consider the risk of problematic unexpected generalization; in other words, concerns about undesirable behaviors despite the system’s formal correctness), security and control (conjugate the selection of the best actions to perform a task and human control over the AI)
- Gain public trust, improve transparency and inform people about how AI works and can be deployed to benefit society at large; one of the biggest impediment to this is represented by black box algorithms, that while on the one hand, they offer greater accuracy than white box algorithms, on the other they stir perplexity and insecurity because their *modus operandi* cannot be fully explained even by its developers.
- Optimizing AI’s economic impact, considering crucial aspects as the labor market, impact on productivity, markets disruption, policy implications, consumer rights.
- Law and ethics of research: the pervasive use of smart machines cannot be understood as free from any ethical and legal restrictions. it is necessary to implement measures such as liability for machines, machine ethics, regulations concerning autonomous weapons and defense systems, data privacy, professional ethics for machines performing a professional work (e.g., for artificial lawyers or artificial surgeons).

International organizations can help scholars and practitioners in defining the path for the future of AI. For example, the International Organization for Standardization created, in 2017, a technical committee about Artificial Intelligence, that goes under the name of ISO/IEC JTC 1/SC 42, whose scope is to “Serve as the focus and proponent for JTC 1’s standardization program on Artificial Intelligence” and to “Serve as the focus and proponent for JTC 1’s standardization program on Artificial Intelligence”<sup>3</sup> (International Organization for Standardization, 2017).

In September 2016, a group of AI researchers stewarding six of the world’s largest technology companies (Apple, Amazon, DeepMind and Google, Facebook, IBM, and Microsoft) announced the “Partnership on AI to benefit people and society” (PAI), a multi-stakeholder organization that aims to formulate and spread best practices on AI-powered technologies, spread knowledge about their potential. Their work rests on six pillars that cover very different topics: security, accountability, economy, system interoperability, societal influences, and social good. To date, it embodies more than 70 partners from 10 countries (and many of those are Non-Profit Organizations) (Partnership on AI, 2017). Compared to the ISO committee, however, the IAP was founded by companies with apparent conflicts of interest, as they are those who can benefit more than others from AI.

The United Nations also contributes to the international debate thanks to a platform called “AI for Good”, launched in 2017 and managed by the International Telecommunication Union. The UN believes that AI is an essential technology for achieving the Sustainable Development Goals (SDGs). The AI for Good is mainly centered on yearly events that aim to bring forward research topics that contribute towards more global problems. The platform also provides an open repository that contains a series of projects, research initiatives or think-tanks that have been considered useful for reaching the yearly SDGs; this enables anyone to connect with other AI stakeholders worldwide so that projects can be more easily developed, for the benefit of society.

---

<sup>3</sup>ISO/IEC JTC 1 is ISO’s committee of Information Technology

## 1.2 The pervasiveness of AI, domain by domain

Alongside big companies, many startups have focused their efforts toward perfecting AI applications for both general purpose and small solutions. This section comprises several real-life applications of AI and ML, to prove the reader how these technologies already produce outcomes in the real-life world. Countless other applications, which can not be conceived today, will come to light during the years to come.

A prominent example of Artificial Intelligence is IBM Watson, a semi-general AI that made its first public appearance in 2011 when competed in a special episode of the American TV show *Jeopardy!*, defeating the human quiz champions by answering real-time questions asked by the host. Nowadays, it is trained to work in several fields, including medicine, finance, and fashion, where it assists humans in decision making.

In China, a high school site in Hangzhou is reportedly testing an AI-based solution to monitor students' expressions and movements for class performance analysis (e.g., if students are paying attention) and improvement (England, 2018). By recognizing emotions and movements, the system provides feedback to the teacher that can act accordingly.

Researchers from Stanford University explored the creative arts by training deep neural networks to generate memes<sup>4</sup>, in an attempt to train it to mime human humor (Peirson et al., 2018); results show that some memes cannot be distinguished (using human evaluation) from those created by humans. Another research group from Stanford University trained a neural network to make predictions about patients using electronic health records (Rajkomar et al., 2018); the model outperformed traditional assessment methods currently used in medicine: the system is capable of predicting the death of a hospitalized patient with an accuracy of 95%.

While a research group trained an algorithm to, by scanning pictures, make assumption about the sub-cultural urban tribe to which a person belongs to (e.g., hipster, punk or surfers), in order to offer her a more personally tailored experience in some services (Kwak et al., 2013), a more controversial application of Artificial Intelligence aims at detecting a person sexual orientation by analyzing her facial images (Wang and Kosinski, 2017). By providing five images, the system has an accuracy of 91% for men and 83% for women; even in this case, the AI is capable of outperforming human judgment by far. Even though this represents an outstanding example of the power of AI, it also raised serious ethical concerns about how it could be used to foster discrimination. University College London Hospitals recently unveiled a plan to implement AI to aid doctors in cancer detection and to decide which patient fast track in the emergency department (Devlin, 2018).

The firm Darwin Geo-Pricing has developed an artificial neural network to perform exploratory pricing. It retrieves the online shopper's location and combines this with data mining to adjust the prices each customer is offered (Darwin Geo-Pricing, 2018). Darwin can determine the optimal price level to pitch at each customer, therefore helping other firms to maximize their profits and compete better with local retailers. By using deep learning algorithms and machine vision, RapidMathematix aims at reducing food waste by applying a dynamic pricing scheme that fixes the price according to the freshness of the product or the closeness to the expiring date. This strategy is, of course, also an excellent way to increase sales, since unfresh products are more likely to be sold (even though at a lower price) (RapidMathematix, 2018).

The platform Mya uses AI to automate some of the tasks involved in the recruiting process, easing it for both job-seekers and recruiters by scaling and speeding up the entire process. It uses Natural Language Understanding to analyze questions and answers given by a candidate and is also able to shift the direction of the conversation to determine how to proceed (Mya, 2018).

Amper is “an artificial intelligence composer, performer, and producer that empowers you to instantly create and customize original music for your content” (Amper, 2018). By easily selecting

---

<sup>4</sup>An internet meme is media content, usually an image or a video, that is spread from person to person within the internet for humorous purposes. They usually relate to a particular internet subculture.

a style, a mood and a length for the soundtrack, the AI can render music that can be downloaded as is or edited according to the user's needs and preferences. Amper does not replace human creativity; it re-imagines and automates the implementation process so that humans can focus on their vision, not the minutiae of production.

General Electric used its proprietary technology in AI to develop digital twins of its products. The digital twin is a virtual duplicate of a digital asset (therefore a bridge between the physical and the digital world) that can be used as “single source of truth for all information related to an asset, including data about past and present state, condition, and performance” (General Electric, 2018). This technology aims at providing more efficient predictive maintenance and at building a digital model of an entire manufacturing plant in order to optimize it at best using real-time data.

The University of Manchester proved that AI could also be used to generate new scientific knowledge. Its AI, named “Eve”, is capable of “develop and test hypotheses to explain observations, run experiments using laboratory robotics, interpret the results to amend their hypotheses, and then repeat the cycle, automating high-throughput hypothesis-led research.” Its main discovery regards how an anti-cancer compound could be used in fighting malaria (University of Cambridge, 2015).

JPMorgan Chase & Co., the biggest bank in the USA, implemented the AI Contract Intelligence (COIN), that can perform just a few seconds document reviews tasks that used to take legal aides about 360,000 hours. Besides the gains in efficiency, thanks to the automation the bank drastically reduced the number of loan-servicing mistakes too (Galeon and Houser, 2017). However, finance is not an isolated case; the industry of law firms is being disrupted by solutions like the one offered by Neota Logic, that enables law firms to automate all the work once executed by paralegals, by summarizing vast amounts of documents, cases, filings and client's data in just a few seconds. Again, the gaining here is not only concerning efficiency, since such a solution is also capable of providing additional information that can strengthen a lawyer's argument (Neota Logic, 2018).

The last example provided is about the food industry; in Chile, The Not Company (NotCo), after realizing how inefficient the R&D process was in the food industry, developed “Giuseppe”, an AI that is capable of generating known food formulas using just plant ingredients instead of animal-based raw materials. It works by analyzing the properties of several plants at a molecular level. This tech company boasts, in its catalog, products like mayonnaise, cheese, yogurt (Shieber, 2018).

## 1.3 The relevance of AI in Economics

The examples from the previous section help to clarify that no sector can ignore the real benefits that the AI can bring to the production of goods, to the provision of services, to the simplification or optimization decision-making processes.

In transportation, autonomous vehicles are the most iconic example, but AI is also used to create complex models of traffic flows in cities. With home/service robots, companies provide new means to deliver parcels or to keep houses clean. AI has disrupted health care since now machines are capable of making predictions and detect disease with greater accuracy than doctors and physicists. Education could be disrupted as well, thanks to personalization at scale of educational programs and a better understanding of the human cognitive processes involved in learning. For low resource communities, AIs can be used to address problems and provide them mitigation or solutions; in developing countries, where statistical data are missing or outdated, AI can be a powerful tool to estimate poverty remotely, assess changes in wealth and poverty, program monitoring and impact evaluation (Blumenstock, 2016). For public safety and security, AIs are extensively used for law enforcement or border security, but also in defense systems like Unmanned Vehicles. In the case of employment and workplace, AI could develop models to redistribute workforce across different geographies or different industries, to reduce the unemployment and

better match one's abilities to necessities. Finally, AI can also be used for entertainment: starting from algorithms that match one's preferences to available products and arriving at machines capable of creating 3D scenes starting from natural language text.

All of these innovations are expected to affect the behavior and the interaction of economic agents and how economies evolve at large. The impact of AI on the economy can be formalized by distinguishing the impact at the macroeconomic level from the impact at the microeconomic level. Both these impacts ultimately affect the position taken by governments and policymakers, whose task is to evaluate the consequences that AI has at different levels (either on the single agent or the whole society) and then promote adequate welfare increasing policies.

Under a macroeconomic lens, AI can either be seen as a potential solution to the problem of stagnant productivity, which has been showing itself for some years now in the most developed economies or as a completely-new productivity factor, alongside the traditional ones. Smart machines are a capital-labor hybrid that is, at the same time, capable of performing labor at scale and requiring capital to be implemented, but also capable of learning from data and act accordingly. Since data, more than money, is what a firm needs to put an AI at work, it is hard to classify AI as pure capital. However, AI is relevant because of its possible impact on occupation and welfare: just as mechanical muscles made human labor less in demand, so are mechanical minds making human brain labor less in demand. On this theme, conflicting philosophies have been proposed. According to the most futuristic-looking scholars, it will indeed be possible to build a general AI capable of out-competing human labor under any aspect. Then, society will have to deal with what has been defined economic singularity: an economy of radical abundance in which no one will need to work anymore, and anyone will benefit from the wealth produced by an unreachable super-intelligence. In a less drastic scenario, AI is still expected to replace humans in certain jobs, since smart algorithms are capable of decision making better and on a scale; this would raise concerns especially about unemployment, wealth distribution and how to retrain and form people to work on other industries. In a final, more realistic scenario, AI could be used for work augmentation, providing humans insights, advice and guidance to increase the firm's productivity, creating more economic value; it is plausible that for many years to come, humans will still be better at thinking outside the box and in elaborate forms of communications. The reader is warned that all of these scenarios are not forecasts, instead, a rhetorical device that aids to frame what the future could potentially look like.

About microeconomics, AI is relevant because it positively impacts the market mechanisms related to the allocation of resources (including data) and the setting of a specific value to the assets in question. Questions about the existence of new market failures arise and, eventually, the goal is how to remedy them. In the case of competition among firms, there is interest in understanding if the AI modifies the current competition models (or if it leads to the definition of new ones), whether this happens between firms that are equally equipped with AI, or whether only some of the competing firms benefit from this technology. In the case of interaction between firm and consumers, the study of AI is vital because of the power that machine learning and predictions can have over the decision making processes of the consumer. If it is true that these features make consumer choices more accessible, more practical and more efficient, it is equally valid that their over-pervasiveness can be detrimental for the power held by the consumer.

Henceforth, the discussion departs from a technological to an economic dimension; from the second chapter onwards, the analysis will almost entirely focus on microeconomics and normative economics.

## 1.4 Macroeconomic impact of Artificial Intelligence

The discussion on the macroeconomic issues related to AI also involved various organizations, each of which offers different perspectives on the impact that the AI will have in the coming decades.

McKinsey believes that automation could help serve as a new productivity engine for the global economy and that it will increase economic growth by 0.8-1.4% in the next 50 years. The



main impact channel will, therefore, be the labor substitution (Manyika et al., 2017). According to their estimation, AI could create annually between \$3.5 and \$5.8 trillion in the global economy (Chui et al., 2018).

PwC, on the other hand, states that the global GDP will increase by around 14% by 2030. This gain will be due to an increase in both productivity and consumption, that will account for 6.7% and 7.9% respectively (Gillham et al., 2018). The contribution to the global economy is expected to be up to \$15.7 trillion in the 2030 economy, \$9.1 trillion of which deriving from consumption side effects, arising from increased customer demand of personalized, higher quality products (Rao and Verweij, 2017). North America and China stand to see the most prominent economic gains in percentage terms from AI: in 2030, it will respectively enhance GDP by 14.5% for the former and by 26.1% for the latter, equivalent to a total of \$10.7 trillion and almost the 70% of the expected global impact (Gillham et al., 2018).

Accenture proposes AI is an entirely new factor of production, not a driver of TFP, and that it could double growth rates by 2035 (for example, US GDP will be 35% higher in 2035). On average, the labor productivity will be 11% to 37% higher in 2035, depending on country (Purdy and Daugherty, 2016). The Gross Value Added (GVA) growth rate is foregone to increase; specifically, in 2035 that of the US will go from 2.6% to 4.6%, that of the UK will increase from 2.5% to 3.9% and that of Japan from 0.8% to 2.7%.

The Economist Intelligence Unit hypothesized three different scenarios of the GDP growth that are likely to happen by 2030. For the first one, they assumed a high degree of complementarity between AI and human skills with investments from the public sector. For the second one, they assumed higher investments from the public sector, especially for granting access to open source data and tax credits. For the third, the most negative one, AI is seen as a substitute for human labor and the public sector shows poor interest in favoring the diffusion of AI. It then compared to a baseline, that stands for The Economist’s current forecast to 2030 (The Economist Intelligence Unit, 2017). According to the analysis, as reported in table 1.1, both the first and the second scenario are favorable, since the estimations for GDP changes are higher in any of the considered countries; the second is preferable since the better attitude of the public sector towards AI fosters growth and favors the adoption of such technology. The third scenario, on the other hand, leads to worse results if compared to the baseline; the losses would be higher than 1% in the considered countries, even leading to negative growth in UK and Australia.

	Baseline	Scenario 1	Scenario 2	Scenario 3
<b>USA</b>	1.84 %	2.04 %	3.00 %	0.84 %
<b>UK</b>	0.63 %	1.29 %	1.94 %	- 1.20 %
<b>Australia</b>	1.03 %	3.11 %	3.74 %	- 0.24 %
<b>Japan</b>	1.57 %	1.96 %	2.43 %	0.53 %
<b>South Korea</b>	1.78 %	2.07 %	3.00 %	0.02 %
<b>Developing Asia</b>	4.34 %	5.04 %	6.47 %	3.20 %

Table 1.1. GDP changes by country by 2030. Depiction of The Economist Intelligence Unit (2017)

### 1.4.1 A regional breakdown

In the race to become the global leader in AI, 24 Countries <sup>5</sup> released strategies to promote the use and development of AI as a major tool to enhance national competitiveness and guard

<sup>5</sup>As of 04th December 2018



national security. No two strategies are alike, with each one focusing on a different aspect of AI policy (Dutton, 2018). Additionally, just like previous technologies did in the past, this one is unlikely to impact different regions in the same way. Its implementation will start in countries equipped with the proper technological infrastructure and with the biggest capability to invest in such technologies and only later is expected to be widespread in poorer regions too. Winners and losers of this race depend on (Baweja et al., 2016):

- Capital for labor substitution,
- Skills and social inequality,
- Technological infrastructures and inertia to innovation, and
- Robustness and flexibility of the legal system.

In the Global Competitiveness Report 2017-2018, the World Economic Forum ranks Countries according to their innovation environment and their technological readiness, using a scale of 1 to 7. As depicted in figure 1.11, the countries with the highest scores are those located in Europe, North America and the Pacific Area: USA, UK, Germany, and Japan before everyone else. Among the followers, China and Italy earned substantially lower scores.




Figure 1.11. Innovation environment (on the left) and technological readiness (on the right) in large advanced economies and large emerging economies, 2009-2017 (Schwab and Sala-i-Martin, 2017)

According to the same source, these five countries are respectively ranked as third, twentieth, fifteenth, twenty-third for the quality of their higher education and training. This ranking takes account for, among everything else, of the quality of the math and science education, both necessary fields for the growth of AI as a new technology. Italy (forty-first) and China (forty-seventh) follow right after.

## China

In July 2017 China's State Council issued a paper, named "A Next Generation Artificial Intelligence Development Plan", which states the guidelines that the country must follow so that it can establish itself as the AI world leader by 2030 (Metz, 2018). In an attempt to compete for the role currently held by the USA (Churchill, 2018), its goal is to create, by 2030, an industry that is worth \$150 billion (China's State Council, 2017). Even if the State Council acknowledges its shortcomings concerning discoveries and inventions of great international impact, as early as 2014, the country overtook the USA regarding the number of research publications it produced - and the number of those that were cited - about the topic of deep learning (Churchill, 2018). Alongside the abundance of capital that the Chinese public sector is willing to invest, China has

the advantage of having the second vital resource for the creation of effective AI solutions: data. Taking advantage of the least barriers (compared to Western Countries) that are opposed to the collection and processing of data, China is amassing huge databases that don't exist in other countries (Knight, 2017). This achievement is partly encouraged by the higher propensity (due to cultural factors) of individuals to share their information. The government is also planning to implement particular policies to overcome the most significant obstacle towards its goal, namely the lack of talents in AI if compared to the USA (Churchill, 2018). Finally, in the paper questions are asked about how to regulate, develop and use the AI ethically.

## United States

In October 2016, the Executive Office of the President National Science and Technology Council Committee on Technology issued a report to provide technical and policy advocacy on subjects related to AI, and to monitor the development of intelligent technologies across industries, the research community, and the Federal Government. The United States recognized the implication of researching about AI, especially in defense systems. In its budget for AI, the Pentagon spent almost \$7.4 billion, an increase of 32% from the \$5.6 billion spent in 2012. However, the public sector is struggling in establishing partnerships with the private sector because of concerns about how AI could be deployed to harm people (Barnes and Chin, 2018). In 2015 the private sector invested \$2.4 billion on AI (CB Insights, 2016), as compared to the approximately \$200 million invested by the National Science Foundation (NSF) (Directionate for Computer and Information Science and Engineering, 2017). The Trump administration's budget for 2018, however, aims to cut science and technology research funding across the government by 10 percent, to about \$175 million (Mozur and Markoff, 2017); this is very likely to result in a shift of R&D to private American companies.

## Europe

Overall, Europe is not keeping the pace of the investment if compared to the Asia and North America: while the first one invested \$3 to \$4 billion in 2016, the second and the third invested respectively \$8 to \$12 billion and \$15 to \$23 billion (Manyika, 2017). On the other hand, Europe hosts the biggest number of the 100 most important research centers about AI in the entire world, 32, compared to 30 located in the USA and 15 in China <sup>6</sup> (Atomico, 2017). The main plan of EU includes more public-private investments, a call for private to share their data to feed machine learning algorithms (while respecting the laws about privacy and data sharing) and social programs to make the transition for automatable works easier for human employees, as well as promote talents in the fields of AI (The European Commission, 2018). The European Commission also calls for the creation of ethics orientation for AI and plans aimed at sensitizing consumers about the effects of the automated decision making. On the wave of European directives, each country is proposing its action plan.

In March 2018, France presented the "AI for humanity" program, a public initiative with which France intends to establish itself as one of the leading countries in the field of artificial intelligence. The strategy presented includes, inter alia, the creation of a European open data framework, guidelines to make it easier for academics to work also for the private sector (up to 50% of their time), the creation of its own technological infrastructure to support the various initiatives, create operational centers to anticipate shifts in the labor market and spread the AI culture on various levels of education. \$ 1.85 billion will be invested in AI projects, divided between public research and startups (Villani, 2018).

The United Kingdom presented its strategy in 2017, in which it considers the relations between government, academy, and industry to be fundamental for the UK to prosper in the field of AI.

---

<sup>6</sup>The ranking considers the number of citations of publications related to AI.

Additionally, in this case, the strong points provide for the creation of data trusts and academic training programs (including over 1.000 more Ph.D. positions in AI by 2025), as well as mixed public-private investment programs (Hall and Presenti, 2017).

Italy, which in 2016 ranked 5th in the world in the production of scientific articles most cited on machine learning after the USA, China, India, and Great Britain (OECD, 2017), has presented guidelines for the public sector so that this can serve citizens and organizations using the AI. Among the fixed points of the program, systems emerge for the management of the educational path and of the working career, environment, tax system (Agenzia per l'Italia Digitale, 2018).

## Middle East

According to their broad definition of AI, PwC estimates the potential economic impact of these technologies in the Middle East to be around \$320 billion by 2030, which is equivalent to almost 11% of the GDP of the region and 2% of the total global benefit of AI in 2030. In terms of geographical shares, the gains are likely to be divided as following: Saudi Arabia (\$135.2 billion), UAE (\$96 billion), GCC4 (\$45.9 billion) and Egypt (\$42.7 billion). Regarding sectors, the biggest gains could be in financial services and healthcare and education (from the public sector) (Jain, 2018). The volatility of the oil price has been, in the past years, a big concern for the Middle East, making it necessary for the countries in this region to seek alternative sources for revenue and growth. By investing public funds, it is important for the governments to set up an environment that fosters the shift towards a highly-technological economy, in which all stakeholders can effectively engage in the development and adoption of AI (UAE Government, 2018; Saudi Arabia Government, 2016). The UAE Government has strategically planned to elevate Dubai into a global platform for knowledge-based, sustainable and innovation-focused businesses, and to provide Dubai with an Autonomous Transportation system, aimed at serving the 25% of the transport demand of the city.

### 1.4.2 Impact of AI on productivity

The macroeconomic study of AI encompasses its likely effect on long-term growth. Every epoch is characterized by a distinct economic problem, and that of the last few years is tied to the sustainable growth of productivity. The question to be asked is whether AI could provide a long-term remedy for such a problem.

Even though, as shown in figure 1.12, the world GDP has grown almost uninterruptedly in the last three centuries, the real GDP growth rates of some individual Western countries have come to an abrupt halt in recent decades, as shown in figure 1.13. Except for China and South Korea, in the period 2010-2017 none of the countries in question showed growth rates above 3%; this result contrasts well, for example, with the five-year period 1985-1989, during which 10 out of 11 countries had growth rates of over 3%.

Besides, as shown in table 1.2, except for South Korea, the average annual Total Factor Productivity (TFP) growth rate in the past decade dropped below the 1%, indeed reaching a negative value in Italy. By then considering the output from computer capital, it can be observed that it has been rising in all the economies portrayed in figure 1.14, even though the rate of the growth declined sensitively, dropping below the 5% for all the considered Countries, except for South Korea and Australia (The Conference Board, 2017). This phenomenon is defined by many as productivity paradox: even though anyone in the developed world has in its smartphone a computing power thousands of times higher than that of computers that allowed Apollo 13 to land on the moon, productivity has reached a plateau from which some fluctuations are possible, but only very slight ones.

Economists have tried to address the slowdown in economic growth in many ways (Mankiw et al., 2011): measurement problems (it is assumed that the data collected do not reflect an actual slowdown in productivity, rather a vice of data themselves), a deterioration of the quality of the




Figure 1.12. World GDP from 1700 to 2016; The vertical lines denote the starting point of each industrial revolution. Source: *The World Bank*




Figure 1.13. Real GDP (constant 2010 US\$) growth rate, by five-year periods and by Country. Source: *The World Bank*

workforce and education, or even the lack of production ideas. According an OECD research, the slow growth of the productivity of the “average” company is actually disguising another situation, namely that the gap between the frontier companies (those that strongly innovate and therefore increase their productivity) and laggard companies (those that take longer to adopt new technologies) is increasingly spreading, especially in the service sector, where the employment of information and ICT is more intensive than manufacturing ([Andrews et al., 2016](#)).

Regardless of the interpretation that the reader gives to the phenomenon, it is in this context that artificial intelligence must be positioned. AI can be defined as a General Purpose Technology (GPT), a technology characterized by pervasiveness, great potential for technical improvement and which favors the birth of the “innovation complementarities” ([Bresnahan and Trajtenberg, 1995](#)). The value of AI does not dwell either in itself or in its theoretical foundation, rather in the versatility that allows people to create innovations based upon it and, therefore, exploiting the wide spectrum of cases in which it is applicable. This is how it finally results in being a major growth driver. Starting from these considerations, we can also evaluate the role of “trailblazers” assumed by the various national governments or by firms such as Facebook, Google or Amazon, who are promoters of the AI upstream (as a research field) and downstream (as commercial applications).

In macroeconomic terms, assuming that the AI systems is likely to result in an increase in

productivity and an increase in the capital stock <sup>7</sup>, modeling the world economy as a closed system that can reach a state of equilibrium, whether if we adopt an exogenous growth model (Solow, 1957) or an endogenous growth model (Mankiw et al., 2011), we would observe an increase in aggregate production and, consequently also an increase in aggregate consumption. A topic that certainly deserves a deepening is how to amortize the “AI-capital”, since although data and IT infrastructures are subject to deterioration, the algorithms may not need to be amortized because its intelligence does not decrease over time.

### Expectations on AI and the Productivity Paradox

Expectations on the impact of AI on productivity are high, and given the so different scenarios proposed by scholars, there is a feeling that discrepancies between the expectations themselves and the statistics detected will occur in the future. These could be attributed to four factors (Brynjolfsson et al., 2017); each one of them is described below.

First, false hopes. There is widespread hype about AI, and people tend to overestimate its realizable potential. This misconception may also be influenced by the common imaginary outlined by literature or the film industry. Nowadays, very little of what Arthur Rodebaugh envisioned during the Golden Age of American Futurism has seen the light, while the rest remains (and perhaps will remain in the future) a pure utopia.

Second, mismeasurement. The point is that the economy is indeed capable of producing more and more efficiently, yet statisticians are not capable of taking full account of the value created. For instance, some people argue that the shortcomings of official statistics are due to the exclusion of free digital products from the GDP computation, as it is widely believed that they are conceptually a non-market (Hatzius et al., 2016). In the case of AI, the problem lies mainly in the intangible nature of the AI output, which makes it easier to overestimate or underestimate its real output.

Third, rent dissipation. While companies try to gain market share or increase their profit from AI-enabled products or services, they end up by dissipate most of the created value while competing among themselves.

Last, lags between the moment when the new technology becomes widespread and the one in which it is found in economic indicators. It has been argued that AI is a GPT, and this means that the more pervasive it is, the more time is needed to accumulate capital stock and to pervade all the industries deeply. For instance, even if there is already the technology for cars to drive autonomously, time is still needed to wait for the autonomous car to become the new standard and, consequently, the measurement lag enlarges. If this logic is applied to every field in which AI can be deployed, it can be seen how the economy needs time before it is possible to assess its effect on productivity.

#### 1.4.3 Impact of AI on employment

A further implication of AI from a macroeconomic perspective is how this technology is expected to affect employment. Since the first industrial revolution, by making human work redundant and less competitive than capital, technologies result in a shift in the capital/labor mix and a change in the composition of labor demand across industries or different geographies. What has been said is already observable among some of the most popular companies that heavily rely on AI already. It turns out that they employ a tiny number of workers if compared to their very high market capitalizations (among the highest in the entire world); it is curious to see that all of them are providing platform services in digital markets.

---

<sup>7</sup>for the purpose of this document, the definition of capital extends so that it can also include AI

	1985-1989	1990-1999	2000-2007	2008-2017
<b>Australia</b>	2,34 %	2,14 %	1,88 %	0,88 %
<b>Canada</b>	2,25 %	1,27 %	1,84 %	0,57 %
<b>France</b>	2,57 %	1,57 %	1,39 %	0,23 %
<b>Germany</b>	2,50 %	1,86 %	1,67 %	0,99 %
<b>Italy</b>	3,25 %	1,44 %	1,08 %	- 0,83 %
<b>Japan</b>	4,49 %	1,24 %	1,33 %	0,64 %
<b>South Korea</b>	8,99 %	6,13 %	4,84 %	2,54 %
<b>Netherlands</b>	2,44 %	2,66 %	1,82 %	0,43 %
<b>United Kingdom</b>	3,94 %	2,09 %	2,28 %	0,35 %
<b>United States</b>	2,88 %	1,99 %	1,67 %	0,64 %


Table 1.2. Average annual TFP growth in selected countries, 1985-2017. Source: *OECD*

Figure 1.14. Growth of Capital Services provided by ICT assets, selected Countries, 1990-2016 (The Conference Board, 2017)

Distributional considerations have emerged as one of the most pressing challenges for policymaking on competitiveness and growth: polarization of wealth and unemployment will be relevant as long as people do not retrain themselves to work in industries that will not be affected by automation. It is also plausible to experience on-shoring trends since the technological infrastructures required for AI to work correctly will be, in the beginning, only available in wealthier countries; the trend will then reverse after that even least developed countries will have a proper infrastructure (Baweja et al., 2016).

According to an estimation, 326 million jobs will be impacted by AI in 2030 (this figure includes jobs which have either been created by AI, AI-dependent or heavily impacted by AI), rather than net jobs created by AI (Gillham et al., 2018). Of course, how the different economies will react to the large-scale adoption of AI for automation depends on the flexibility of their labor market. In this sense, it is useful to look at the index proposed by the World Economic Forum depicted in figure 1.15. By ranging on a scale 1 to 7, computed using many parameters (business executives' perceptions of union-employer cooperation, flexible hiring and firing practices, and the alignment between wages and productivity), shows that in the past ten years the labor market flexibility is slightly decreased in every region, except for Europe (where higher flexibility is experienced) and for the Middle-East and North Africa (where the flexibility is relatively stable). According to this

	Market Cap [Billion \$]	As of	Employees	As of
<b>Airbnb</b>	31	March 2017	3.100	March 2017
<b>Spotify</b>	33,71	August 2018	3.969	July 2018
<b>Netflix</b>	150,78	August 2018	5.500	December 2017
<b>Facebook</b>	512,46	August 2018	30.275	June 2018
<b>Alphabet</b>	863,1	August 2018	89.058	June 2018
<b>Amazon</b>	925,66	August 2018	566.000	December 2017

Table 1.3. Market capitalization VS number of employees; selected companies. Source: *ycharts.com, forbes.com*

statistic, the effect of AI on employment is more likely to be absorbed in the Pacific Ares, Middle East, Europe, and North America.




Figure 1.15. Evolution of labor market flexibility by region (on the left) and within the European Union (on the right), 2007-2017 (Schwab and Sala-i-Martin, 2017)

## The task-based approach to labor market

Unlike other forms of capital that have in the past replaced work, the AI requires a fresh approach to the study of how much, at the production level, it is replaceable. To study the labor market dynamics in this context, labor (hence, production) is modeled using a task-based approach, in which: (1) the output is created by combining the effect of various tasks and (2) capital and labor are replaceable with a specific substitutability rate (the choice to allocate a task to work or capital depends mainly on the technology, that is, if the AI can perform it or not) (Acemoglu and Restrepo, 2018). As the price of the factors varies, the range of tasks allocated to it and the incentives for the introduction of new tasks vary as well.

An extensive framework proposed by Acemoglu and Restrepo considers both the automation of tasks that were previously executed using labor and the introduction of new tasks in which labor has a comparative advantage over capital (new tasks are created as soon as old tasks are automated, even though these two phenomena has different growth rates). This framework also highlights the price between the production factors and the capital/labor shares. The principal finding of the research is that if the comparative advantage of labor over capital is sustainable and the number of the newly created tasks is sufficiently high, the demand of labor can remain stable (or even grow) over time, despite the process of automation (Acemoglu and Restrepo, 2016). This result, however, means that the demand for labor is addressed to increasingly skilled workers since



the low-skill tasks can be easily automated and be performed by capital. The gap between low- and high-skilled workers may rise, leading to more severe redistribution concerns.

### Jobs destruction

Talking about employment and new technologies, most of the concerns are about the adverse effects on the labor market, like which and how many jobs are going to be depleted. Many researchers and practitioners have then focused their efforts on assessing the impact of automation on the job market. Using a sample of 702 occupations, classified according to the characteristics of the tasks that compose them, it emerged that 47% of US jobs are at high risk of computerization, while 19% are at medium risk. Regarding industries, those most at risk are services, sales, and constructions. The works that are less susceptible to computerization are those that require perception and manipulation, creative intelligence or social intelligence (i.e., those that, at the current stage of ML development, are not yet wholly automatable) (Frey and Osborne, 2017).

McKinsey, in the other hand, decomposed 800 occupations in 2000 simpler tasks. It emerged that, at the current state of technology, only 5% of all occupations can be entirely automated, and that around 60% of occupations have at least a 30% component that could be automated. The activities with the highest automation potential are those performed in highly structured and predictable environments, namely: those involving predictable physical activities (e.g., warehousing workers) for the 81%, the data processors for the 69% and the data collectors, for the 64% (Manyika et al., 2017).

PwC proposes a different investigation, in which Countries are grouped in regions and jobs (rather than the tasks that constitute them) are organized by industry. Generally, North America and Europe have the highest rates of jobs at risk of automation; concerning the industrial sector, Transport and Logistics, Energy & Utilities and Manufacturing are those for which higher automation rates are estimated (Gillham et al., 2018).




Figure 1.16. Percentage of jobs at risk of automation, by 2030, by geographical region and industry (Gillham et al., 2018)

### Jobs creation

A good aspect about such a transformative technology is that it can also positively impact the employment landscape. In such a rapidly evolving scenario, an estimate states that 65% of children entering primary school today (meaning, those born between 2011 and 2013) are likely to end up



working in jobs that do not even exist yet (World Economic Forum, 2016). To prevent a worst-case scenario (technological transformation followed by talent shortages, mass unemployment, and increasing inequality) reskilling and upskilling of today’s workers will be critical.

Even if the future scenario is highly unpredictable and therefore difficult to picture, it is still possible to make short-term considerations. For instance, there is a growing demand for highly-skilled individuals able to create value by developing (or working with) the new technology. As depicted in figure 1.17, since 2013 the share of jobs requiring AI skills on the web portal Indeed.com is around 4.5, 8 and 12 times the share of 2013 for USA, UK, and Canada respectively. By considering the job openings on the portal Monster.com, it is possible to see how Machine Learning, Deep Learning, and Natural Language Processing are the most requested skills (Shoham et al., 2017).




Figure 1.17. Share of jobs requiring AI skills on the portal Indeed.com (selected Countries), by year (Shoham et al., 2017)




Figure 1.18. Job Openings, disaggregated by required skills on the portal Monster.com, by year (Shoham et al., 2017)

The great pervasiveness of AI will also lead to the creation of new AI-driven business and technology jobs, that can be grouped in three main categories (Wilson et al., 2017):

**Trainers** People that will teach AI technologies how to perform and where possible to mimic human behaviors for chat-bots or virtual assistants, including how to show compassion, detect sarcasm and use humor in appropriate situations.

**Explainers** Professionals that “bridge the gap between technologies and business leaders”, technicians who can explain how algorithms (especially black boxes) work and understand why the response and output is a certain conclusion. They can also be employed to determine which specific algorithm has to be used for a specific task.

**Sustainers** Individuals who evaluate the non-economic aspects of AI, such as ethics, to resolve the unintended consequences that could arise by the use of smart algorithms.

Kai-Fu Lee, a precursor in the field of speech recognition and AI expert, argues that jobs involving repetitive, routine or optimization tasks (e.g., customer support, hematology, reporting) are the ones most at risk of being replaced by intelligent machines. On the other hand, works with greater creative or strategic content (scientist or economist) for the coming decades are far from being replaced (see figure 1.19). However, this does not rule out that AI cannot assist people even

in more creative work or in those where empathy and human feelings play a central role (Kai-Fu, 2018).

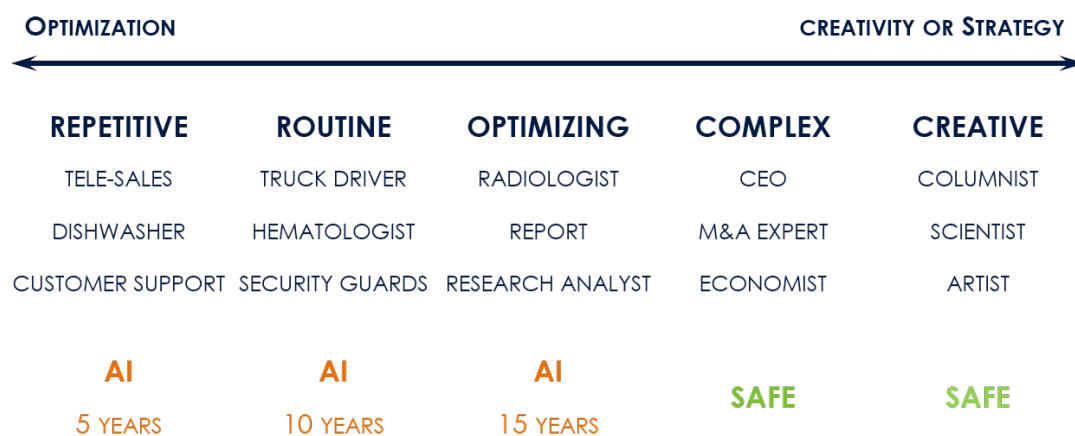


Figure 1.19. Classification of jobs according to the Axis of creativity, as proposed by Kai-fu Lee.

## 1.5 Policy implications of AI

Powerful technologies can produce significant benefits, but they can often produce great harm; AI makes no exception. Regulators and companies should be aware of their potential to cooperate and define the best possible path for the future; given the technical knowledge necessary to understand AI, the governance is expected to involve more experts who can understand and shape interactions between society and smart machines so that no one is left behind. It is equally critical to raise awareness about ethical, privacy and security issues that could arise. Although Europe is already on the right path thanks to the recently-approved GDPR, the same cannot be stated for other parts of the world, where the privacy laws are much less restrictive or, in some cases, wholly absent or not applied.

This section, which is everything but an exhaustive argumentation, addresses some of the hardest challenges for the regulator; all of them embody some form of ethical issues and therefore require robust diligence and commitment to be adequately faced. The topics about data privacy and “algorithmic fairness” will be better detailed in the next chapters of this work.

Perhaps the most critical problem is related to the privacy of those using AI-based products or those whose data are used to put the AI at work. Many people are understandably concerned about their information being misused by organizations, corporations, governments, pressure groups or even individuals. This is a consequence of not providing enough transparency to customers, due to organization’s secrecy on how data are used and difficulty in explaining how the ML operates a prediction: unless individuals are equipped with proper knowledge and control, they will be subject to decisions that they do not understand and have no control over. It is therefore essential to address who is the data controller for an autonomous machine with self-learning capabilities, ensure that the data managers adopt proper countermeasures to prevent data to be stolen, misused or sold improperly, make it impossible for organizations to collect from its users more data than strictly necessary.

About the ethical dimension, AI could spark two problems: objectification (concerning violation of human dignity) and stigmatization (concerning using AI to predict human behavior). This led someone to question if it is necessary to have privacy engineers to embed privacy by design features in novel products (The European Commission, 2016). The biggest concerns are about AI perpetuating society’s biased discrimination based on faith, race, sexual orientation, social rating

systems, that could be enforced (perhaps accidentally) by AI. This class of problems is already encountered by developers when designing AI; for instance, people often take risks and break the law in case of need, like when they break the speed limit to get out of danger. Should machines also break the rules in this way and, if so, by how much? Also, again, how to teach a machine when how much is enough?

Algorithms can discriminate, especially when these algorithms learn from data. It is important to consider that the training sets used for Machine Learning could bias the AI: this is possible because the machine has no other data to contrast the information contained in the first training set with the bigger picture. It is, in fact, true that a tech developer who wants to produce a biased algorithm can do so, but in practice, even an unbiased developer with the best intentions can inadvertently produce a system that returns biased outputs.

About, employment, a change in the capital/labor mix also has consequences for the redistribution of well-being; since the economy, due to various market frictions, is not Pareto improving, innovator and workers may not benefit in the same way as automation. In this sense, it is the duty of the institutions to promote compensation mechanisms that allow, in the case of job replacement in favor of capital, to counteract technological unemployment (be they of a monetary nature or training so that they can be used in areas where they are replaceable from the capital) (Korinek and Stiglitz, 2017). Besides, public policy should also address the problem caused by AI doing jobs that usually require certifications (e.g., if a surgeon needs a degree in order to operate patients, what about an AI-surgeon?). Finally, the use of AI-enabled products that make consequential decisions about people, often replacing decisions made by human bureaucratic processes, may raise concerns about how to ensure justice, fairness, and accountability.

Last but not least, scholars question the civil or criminal liability of smart machines. In the event of an accident, it is necessary to understand to what extent responsibility lies with the developer and how much on the machine itself; and if an AI system is held liable, should it be held liable as an innocent agent, an accomplice, or a perpetrator? (Hallevy, 2010) proposes three models that discriminate against cases where AI may be subject to criminal liability from those in which criminal liability is attributable to developers, vendors or end users. Whether AI systems can be held legally liable depends on three factors (Kingston, 2016): (1) the limitations of the system, and whether these are known and/or communicated to the purchaser; (2) the nature of the AI system (product or service); (3) whether the offense requires a mental intent or is a strict liability offense.

## 1.6 Microeconomic dynamics brought by AI

Much of the research carried out by companies is driven by the possibility of creating new markets or disrupt existing ones. To date, the consumer is (more or less consciously) to interact with the AI of different nature who deal with different tasks: automated SM, media industry, VPAs, recommendation engines. Each of these interactions involves many unusual dynamics.

An interesting topic about the economics of AI is how a decision maker can be influenced in its final choice by algorithms, predictions, and responses given by AI itself. The key is to understand how much power can be left to the algorithm and how much to the decision maker. It is crucial for an organization to provide the agent what she wants, not what the AI wants (or wrongly believes she wants): For instance, the best machine-learning algorithms based on patterns favor proximity above diversity, which is all but how humans have evolved.

As already pointed out, one of the most considerable advantages that AI gives to consumers is saving time. According to Gartner (Forni, 2017), in 2017 500 million users were enabled to save up to two hours a day thanks to AI features embedded in every day's products and services. Machines are far better than humans at managing multiple factors at once when making complex choices, can elaborate much more data at once, and apply probability to suggest the best possible outcome. This phenomenon works on two levels: first, a reduction of the search costs associated to the identification of a desired product or service translates into a higher perceived utility

and eventually, in a higher consumption rate (just because it is easier to access to the goods in question). Besides, if the AI frees the consumer from the execution of non-value adding tasks, she has more spare time that (again) could lead to higher consumption rates.

Algorithms can also be used to apply price discrimination, dynamic pricing or customer profiling to create new opportunities for a company to sell its products or create new personalized ones to better match customers' demands. Consumers today expect outstanding personalized experiences that push rather than pull. Predictive analytics allow marketers to target audiences better, reaching them with content they care about.

Data acquired from customer searches and buying behaviors are used to customize content at the individual level, while insights gained through cognitive intelligence drive smart recommendations for tailored experiences that shorten the purchase journey. However, AI is advancing beyond data analysis and rushing into data production, streamlining the content-creation process. Intelligent automation software can help brands create on-demand advertisements, summaries, and articles from structured data. Content that is automatically generated from data inputs makes delivering messages across multiple platforms speedy and precise.

Whether this is seen as sophisticated, micro-targeting marketing or just technology, it is essential to assess the impact that these features could have over the utility (or similarly, on the surplus) gained by the consumer, how much the consumption patterns can be influenced so that a company can sell more than one product or even products decided by her. No less important are the competition dynamics of a company that adopts AI and a company that does not do it, as well as the competition between companies that implement AI in the same way, as these dynamics do not embody an isolated sphere but have (more or less intense) repercussions on the final consumer as well.

The work presented in the following chapters starts with a quantitative model of microeconomic interaction between the consumer and the company, operating simultaneously on two interconnected markets. The findings of the model are then enriched by a qualitative treatment of the implications that the dynamics highlighted by the model have for the consumer and his well-being.



## Chapter 2

# A microeconomic model of customer-firm interaction

The topics discussed so far deliver a powerful message: AI is essential for economics and, because of its tight bond with data, it will result in a change in firms and customers interplay in different markets. As already pointed out, if in the past decades products have been subject to electrification and, later, digitalization, nowadays products and services are subject to smartification through the implementation of AI features. Apart from smartification, AI also enabled utterly new products and services that leverage the power of predictions and pattern recognition to create value.

The main difference from the firms that do not offer smart products is that, for the latter, the quality associated to the products depends solely on the direct effort that the firm exerts in R&D, while smart products continue to improve from the data gathered from their users. Consumers are then an active part in the process of product improvement because their interaction with the firm (that takes place through the interaction with the product) is observed, codified and used to create value. For instance, Netflix's CEO, Reed Hastings, stated several times that the ultimate goal of the company's recommendation algorithm is to suggest to their users a different movie for each of their different current moods. If this will be possible one day, it is evident how a pure streaming service with no recommendation features will find a hard time competing with the former.

### 2.1 ML: an economics perspective

From a technological perspective, ML can be seen as the interaction of three blocks, namely data, algorithm, dynamic evolution. If the goal is to offer an economics-perspective of ML, the three blocks can be transposed using notorious economics tools, as described below and depicted in figure 2.1:

- data are modeled as one side of a multi-sided market (another side is represented by the product market, where the product embeds ML features),
- the algorithm necessary to offer these features can be modeled as an investment that the firm has to commit to, and
- the dynamic evolution of the ML system can be modeled using a multi-stage game.

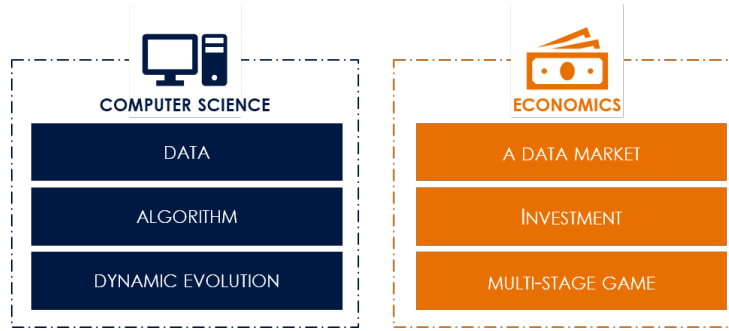


Figure 2.1. ML from a technological and from an economics perspective.

## 2.2 Model scope

The model presented in this chapter has a double scope:

- Starting from the analogy of section 2.1, explain how a firm and its customers interact when the former offers to the latter a product with ML features, and
- Explain the strategy pursued by the firm when it decides to offer an ML-powered product in a certain market.

Even though AI features space in very different domains, it is possible to reduce the problem to three different features, namely (1) the predictions, (2) recommendations, (3) customer perspectives about shared data and future gains. Make a prediction means take the information that an agent has and use it to generate information that the same agent does not have. In many cases, the smartness showed by AI-enabled products is nothing but the ability to make predictions about their users. The firm has to commit an investment in smart algorithms in order to be capable of offering these features. Predictions are the main ingredient for a successful recommendation recipe. By leveraging on recommendations, customers could, for example, benefit from lower search costs and avoid choice overloads, but could also result in excessive persistence of the system, a lower scope in product diversity and finally, in worse customer experience. Without data, algorithms are useless, since the smartness (hence, the value) is built upon the information contained in them. The data that customers share with the smart products can assume three different roles, according to the timing of their usage and collection. The capability of a machine to be smart depends on its capability to gather data from its users. Such trait, in contrast, raises concerns about privacy and clashes with the different perception that people have about sharing personal information with third parties. A depiction of the logic process that, starting from data, leads to predictions, is shown in figure 2.2.




Figure 2.2. From data to prediction: logic process depiction.

## 2.3 Model structure

Hence, the author presents a microeconomic model that attempts to take account of the dynamics described above, in which the firm and the customers interact simultaneously on two interconnected markets: one for the product and one for the data. The model structure is depicted in figure 2.3.



Figure 2.3. Model breakdown structure.

The remainder of this chapter is structured as follows: first, the chapter introduces and explains all the assumptions under which the model is built. A description of both consumer's and firm's behavior follows. It later introduces how the consumers and the firm interact in a generic stage of the game of the defined environment; then, a description of how customer and firm interact dynamically is proposed. Finally, the firm's strategy in the first stage of the game is presented. The last three points are analyzed in two alternative scenarios: one with a monopolist firm and another one with a social optimizer firm. This comparison helps to understand which solution (under which assumptions) is to be preferred to the other one. The chapter ends with a numerical example that consists of a possible application of the model, and a section that summarizes the main findings of the model proposed.

For the reader going through this chapter, it is helpful to keep in mind a few popular services like Spotify, Google Maps, Netflix or any other commercial applications that many people use more and more every day, known for their massive use of ML. Nevertheless, the model presented below appears to be robust even if applied to less known AI-enabled products or services. To highlight the implications of this model, this is compared to the classic microeconomic model that involves a monopolist firm and their consumers.

## 2.4 Two-sided market setup

For this analysis, the problem is set up as follows. A single-product firm (a monopolist first, a social optimizer second) offers its product <sup>1</sup> that embeds certain AI features on a particular market. The product considered for the analysis is a non-durable good, so that consumption occurs in the period of purchase. On the other side, consumers demand a certain amount of the product according to a specific utility function that has to take account of the value of the product demanded, of the value arising from the AI features and of the value of the data that the customer needs to share with the firm.

According to this setting, the firm and the customers are also (and simultaneously) operating on a second market, interrelated to the first, in which data are exchanged. More precisely, because of the interactions that take place on the monopolist's platform (the first market), customers are producing and sharing a variety of data about their tastes, preferences, habits, creating a data platform (the second market) that can be accessed by the firm in order to create value (for both the firm and the customer).

These data are assumed to be an inseparable byproduct of the customer interaction with the firm so that each customer transaction always works on two levels: it represents the sale of the

---

<sup>1</sup>Henceforth, if not otherwise specified, the term product is used to indicate a generic AI-powered good or service offered by the firm.



product and the harvesting of customer data. Data sharing requires both the company and users to trade off. Both agents can benefit from more data (the company can improve its Machine Learning system, and the customers can receive products of higher quality). However, at the same time data are costly to acquire for the firm and costly to give for the consumers (especially concerning the loss of privacy).

The market interconnection makes the data market have a network effect that enables the Machine Learning-powered product to become smarter as it gets more data from its users. As the quality of the first market improves, there will be a higher demand for the product and, because of another network effect, this will translate into a more significant data collection on the secondary market; the situation is depicted in figure 2.4. The smartness of the Machine Learning system is, therefore, a proxy for the quality of the product: the more precise the algorithms, the higher is the quality, allowing the firm to exert some power in pricing its product. The magnitude of the data network effect will affect the strategic decisions of the firm.

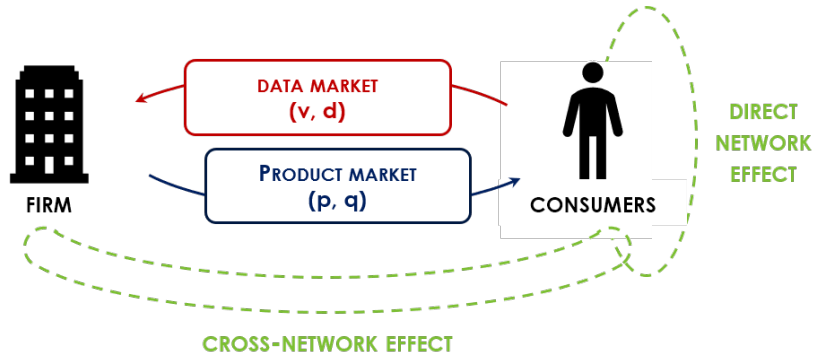


Figure 2.4. Depiction of the consumer-firm interaction, including the direct and cross-network effects.

## 2.5 Consumer's behavior

The first step to set up a model that considers all the factors described so far and that can be used to draw conclusions about the role of the data in the customer-firm interplay, it is essential to define demand functions. Here, consumers are assumed to be rational, meaning that they make reasoned decisions that can provide them with the highest personal utility. The individual linear demand is derived as a solution of the utility maximization problem made up by a linear-quadratic utility function and two budget constraints:

$$U(q_0, q, d) = (V + U_{ML} - c_{BIAS}) \cdot q + r_{ML} \cdot d - 1/2(\alpha q^2 - 2\beta qd + \alpha d^2) + q_0 \quad (2.1)$$

$$s.t. \begin{cases} pq + q_0 \leq M_B \\ vd \leq D_B \end{cases}$$

The equation 2.1 showed above implies that the utility enjoyed by a generic consumer is proportional to the utility of the product itself, plus the additional utility arising from the AI features of the product and minus the disutility due to the bias of the Machine Learning system. The consumer can also benefit from additional utility arising from its data shared with the firm, that could, for example, take the form of more accurate personalization or profiling. Finally, the variable  $q_0$  represents the Hicksian composite commodity, that contains all other goods outside the market under consideration; the price of one unit of this basket is normalized to 1 so that it does not affect the study of the model (Belleflamme and Peitz, 2015).

On the other hand, the two budget constraints mean that (1) the consumer is going to consume the product as long as the budget constraint  $M_B$  allow him to buy an additional unit of the good and (2) that the consumer is willing to share only a limited amount of data with the firm, as imposed by the variable  $D_B$ .

It is assumed that data and the product sold by the firm are complementary products, in a proportion that is fixed by the values assumed by the variables  $\alpha$  and  $\beta$ . The hypothesis of product independence has been excluded since empirical evidence shows how usually the price of a product that requires the user to share some of her data seems to be somehow influenced by the quality or the quantity of the shared data. Finally, the hypothesis of substitutability has been discarded since data cannot replace a product and vice-versa. When derived, the utility function gives rise to the following inverse demand functions:

$$\begin{cases} p(q, d) = V + U_{ML} - c_{BIAS} - \alpha q + \beta d \\ v(q, d) = r_{ML} + \beta q - \alpha d \end{cases} \quad (2.2)$$

The price becomes a strategic variable to attract customers. It depends linearly from the quality of the Machine Learning system (through the net effect of the variables  $U_{ML}$  and  $c_{BIAS}$ ), meaning that, during the first sale periods, the product will result least attractive because of the uncertainty about the added value arising from the data pool, and the price will be lower to ease the purchase. On the other hand, the value of data is increased by the dependency on the demand of the product; the data value is, however, decreasing when the customer shares more data.

This result is straightforward since it is true that if, for instance, a user shares data for one hundred days, the data collected in the first few days will contain more information than the data collected on the last day of the observation. This latter system of equations can be inverted (this is true  $\forall \alpha, \beta | -1 < \beta/\alpha < 1$ ) to obtain direct demand functions. For better readability of the result, let  $\phi = 1/(\alpha^2 - \beta^2)$ , which is always a positive number because of the assumptions about  $\alpha$  and  $\beta$ . The demand functions then take the following form:

$$\begin{cases} q(p, v) = \phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] \\ d(p, v) = \phi[\beta(V + U_{ML} - c_{BIAS} - p) + \alpha(r_{ML} - v)] \end{cases} \quad (2.3)$$

The system of two equations above highlights the double role of buyer of product and seller of data at the same time. Specifically, the function  $q(p, v)$  indicates the quantity of product that the consumer purchases from the firm, while the function  $d(p, v)$  indicates the quantity of data that the consumer is exchanging with the firm while purchasing the product.

The assumption of complementarity holds in both cases since in both equations the prices have a negative sign and are multiplied by positive coefficients only. The purchased quantity is finite even in the case of a null price. The net effect of the Machine Learning system depends on how much this is biased. The net effect of the data on demand (the addend  $((r_{ML} - v))$ ) depends on how much the customer values the current value of its data if compared to the (discounted) utility that he could benefit from by sharing the data.

Similar considerations hold for the amount of data that the customer shares with the firm, but the effect of each variable is influenced by the contrary amount ( $\alpha$  instead of  $\beta$  and vice-versa). In this case, however, even though the consumer assumes the role of a seller, the amount of data shared decreases when the value of the data increases. This result may seem a contradiction, given that in economic theory when the price of good increases, its supply increases. In this case, the rational consumer, on the other hand, is aware that sharing information of very high value (e.g., extremely confidential data) violates his privacy beliefs, and therefore will be less inclined to share.

The general equations derived above are specific to the static customer-firm interaction, that corresponds to specific performances of the ML that, as already pointed out, is an evolving system. The current demand level is also tied to the future level of quality of the product: a higher demand

today means that the firm is collecting a more considerable amount of data that, when added to the current stock present on the platform, will allow the Machine Learning system to become more efficient. The quality level in the future means two things:

- A higher level of individual demand, and
- an increase in the number of customers ( $N_t$ ) because of the network effect assumed

This improvement process only stops (to become constant from there on) when the algorithms have a perfect accuracy: it is, therefore, essential to study the problem in the form of a dynamic game.

### 2.5.1 Consumer-related variables

Apart from the utility that the consumer gains from consuming a unit of the product (denoted as  $V$ ), the consumer is subject to additional effects due to the AI features of the product. Specifically, the consumer can also benefit from an additional variable utility arising from the Machine Learning features of the product; this can both be expressed in function of ML efficiency (equation 2.9) and in function of cumulate data (this is true since the ML efficiency is itself a function of the cumulate data):

$$U_{ML} = \begin{cases} f(\eta) \\ g(D^T) = V \cdot (1 - e^{-\eta_{DATA} D^T}) \end{cases} \quad (2.4)$$

In the first case, the domain of the function  $f(\eta)$  is equal to the image of the overall ML system performances ( $\text{Dom}(g(D^T)) \equiv \text{Im}(\eta) = [0; 1]$ ) while its image is limited to the range  $[0; V]$ . Accordingly, the utility arising from the ML feature is monotonically increasing in the efficiency of the ML system and that it will be maximum when the ML system is perfect. In the second case, the domain of the function  $g(D^T)$  corresponds to the range  $[0; \infty]$ , while its image is still the range  $[0; V]$ <sup>2</sup>; this means that the utility is increasing in the accumulation of the data stock, which can be virtually infinite.

The reason why the variable  $U_{ML}$  depends on the data through a negative exponential function is rather straightforward: the AI start by picking all the low-hanging fruits and therefore improve very quickly (for which quality jumps will be significant and visible), but then it runs into some difficulties. Infusion times, but the magnitude of smartness increase might be smaller and smaller (for which quality leaps will be infinitesimal and negligible).

The utility quantified in 2.4 is distinguished from another value, denoted as  $r_{ML}$ ; through this distinction, the private benefit is kept apart by the social benefit that arises from consumption. Since the proposed context is dynamic, consumers have to hold beliefs about the future course of action;  $r_{ML}$  models the utility that the consumer expects to yield in the future. To illustrate, a consumer shares his data today so that in one year the accuracy of the recommendation will be higher, and she will be able to enjoy greater utility when using the product; at the same time, the consumer is also aware that the decision to share data exposes her to a higher risk of privacy-related issues. The condition  $r_{ML}$  is a strategic variable chosen by the consumer (and not by the firm) since she is the one deciding to share her data with the platform. This model does not take account of the decaying value of data over time, even though this is demonstrated to be an influential phenomenon in real cases.

$$r_{ML} = h(\eta, \delta) \quad (2.5)$$

---

<sup>2</sup>In the proposed model, the UML variable is constrained, by assumption, to the maximum value  $V$ . This assumption is not necessarily verified a priori, since there may be cases in which the value given by the smart features far outweigh the value of the product itself (in the extreme case in which all the value comes from the AI, this could also not exist!)

where

$\delta \in [0; 1)$  is the generic discount factor of the future utility.

Lastly, the quantity introduced below models the effect that a biased algorithm has on the overall utility of the consumer. Bias can be modeled as a deviation from the optimal output of the ML algorithm, hence as a negative externality that makes the benefits arising from the shared data pool decrease. The bias presented in this model is meant to be arising just from an AI that has not been trained with enough data to provide a perfect output (it does not model, for instance, bias used as a strategic variable of the firm). The cost of bias is assumed to be proportional either to the stock of collected data or the ML system efficiency:

$$c_{BIAS} = \begin{cases} w(\eta) \\ z(D^T) = V \cdot e^{-\eta_{DATA} D^T} \end{cases} \quad (2.6)$$

The modeling of the bias cost is specular to the definition of  $U_{ML}$ . Again, in the first case, the domain is equal to the image of the overall ML system performances ( $\text{Dom}(g(D^T)) = \text{Im}(\eta) = [0; 1]$ ) and the image is limited to the range  $[0; V]$ , where  $V$  is the utility associated to the consumption of a unit of product. For this reason, the bias cost is monotonically decreasing in the efficiency of the ML system and that it will be null when the ML system is perfect. In the second case, the domain corresponds to the range  $[0; \infty]$ , while the image of the function is still the range  $[0; V]$ ; this means that the bias cost is decreasing in the accumulation of the data stock, which can be virtually infinite.

### 2.5.2 Discussion of model assumptions

The quantity  $U_{ML} - c_{BIAS}$  incorporates the net effect of the interdependency between users, implying that each customer receives benefits from the service according to the number of users. Even though the dependency is on the cumulate amount of data shared by users, it is assumed that each one of them acts as the representative customer; by induction, the dependency moves to the number of customers.

The situation is better explained using the figure 2.5, where the value  $V$  has been normalized to 1 for representational purposes. Each period is associated with a particular stock of data that, when used to feed the ML system yields a certain efficiency level. This effect, in turn, is associated to a known level of net additional benefit that will be (1) negative before the intersection point (meaning that the cost of bias is higher than the additional utility), (2) null at the intersection of the curves and (3) positive after the intersection. In correspondence of an infinitely big stock of data, the net effect is maximum and equal to  $V$ .

In real applications, the limit case with  $U_{ML} = 0$  and  $c_{BIAS} = V$  is improbable to be found since the AI is trained to give a determined type of output when queried before being used in commercial applications. Indeed, this step is crucial for the firm; otherwise, there would be value destroying, rather than value creating, for the customer.

The net effect of the ML system is crucial in the exploitation of the network effect, since a more efficient system attracts more users that will, in turn, share more data with the platform, leading to even higher efficiency, and so on in a positive feedback loop that will ultimately settle when the ML system is perfect (meaning, for example, perfect and unbiased recommendations for the user).

When this result is achieved, the customers are expected to keep sharing data with the platform since they want to keep benefiting from the high-quality delivered by the ML features.




Figure 2.5. Trend of the variables  $U_{ML}$  and  $c_{BIAS}$  versus the stock of collected data

## 2.6 Firm's behavior

The situation depicted so far implies that the firm has a different solution (in terms of strategic variables) for each different level of ML efficiency. On this wise, in order to quantify the optimal investment level that maximizes the firm's payoff, it has to play a multiple-stage game with the following strategic situation:

- In period 1, the firm maximizes its payoff by choosing the optimal investment in ML algorithms, and
- In all the other periods, the firm fixes the optimal values for the price of the goods sold and for the demand of data that it has to collect from its consumers.

Because of the data-driven indirect network effect, the level of investment is lower than the case in which the firm attempts to build a perfect quality system from the beginning of its business (this would theoretically be infinite since it would take the firm to program every possible scenario of customer-firm interaction). Without considering that it is possible to extract much more knowledge from the data than it is possible to encode in a machine, why pay experts to slowly (and painfully!) encode knowledge into a form that machines can understand when all of this knowledge can also be extracted from data at a fraction of the cost?

### 2.6.1 Firm-related variables

To offer a product that embeds AI features, the generic firm has to commit to an investment, denoted as  $I(\eta_{ALG})$ , which is operation-independent and that affects the profit level of the firm. The size of the investment depends on the quality of the Machine Learning algorithms that the firm wants to implement in its smart system; since the efficiency of a generic algorithm cannot be (by definition) greater than 1.0, the investment needs to be capped to a maximum threshold. Such effect is obtained by using a negative exponential function like the following, which is increasing and convex in  $\eta_{ALG}$ :

$$I(\eta_{ALG}) = A + K[1 - e^{-\lambda \cdot \eta_{ALG}}] \quad (2.7)$$

where,

$A$  is a sunk cost that the firm incurs in if it decides to undertake the investment (e.g., licenses, legal permits, hardware, ...),

$K$  is equal to the total investment capacity of the firm minus the initial investment  $A$ ;  $K$  must satisfy the condition  $(K \geq I(\eta_{ALG}) - A)$ , and

$\lambda$  is a generic positive constant.

Because of the investment, the resulting cost function of the firm can be written as:

$$C_{TOT}(q) = I(\eta_{ALG}) + cq \quad (2.8)$$

where,

$c$  is the constant marginal cost associated with an additional unit of output, and

$q$  is the quantity of product supplied by the firm.

It has been argued that the quality of an AI's decision depends on the quantity and the quality of data; when combined with data, the Machine Learning system shows certain performances that are also a function of the algorithm performances. It is, consequently, possible to define the overall Machine Learning system performances as:

$$\eta = 1 - e^{-\eta_{DATA} \cdot \eta_{ALG} \cdot D^T} \quad (2.9)$$

where,

$\eta_{DATA}$  is an exogenous variable that explains the quality of the data that feed the algorithm (it is, in fact, unrealistic to think that all the data collected by the firm yield the same output or that they are exempt from noise or errors), and

$D^T$  is the cumulate amount of data collected by the firm. The equation 2.9 implies that there are no flaws arising from the collection of too many data; in other words, by assumption, more data are always associated with higher ML performances; this assumption arises from the empirical evidence that the quality of a recommendation engine depends more on the scale of the pooled data rather than on the power of the algorithm (Schaefer et al., 2018a).

Just like people's beliefs are based on their experience, which gives them an anything-but-complete picture of the world, and usually leads them to jump to false conclusions, algorithms owe their intelligence to the experience too, which in their case takes the form of data taken by the user. The function  $D^T$  implies that, because of its effect on the Machine Learning efficiency, as the number of users increases, the benefits that each customer can obtain from the service increases.  $D^T$  is defined as the following summation:

$$D^T = \sum_{t=1}^T N_t \cdot d_t \quad (2.10)$$

where,

$d_t$  is the data shared by the  $i$ -th customer with the firm in the period  $t$ , and

$N_t$  is the cumulate amount of customers of the firm in the period  $t$ , and  $N(t)$  is a subset of the total potential users  $N_{MAX}$ . The trend of  $N_t$  over time can be modeled using an S-shaped curve.

### 2.6.2 Number of sub-games

For each sub-game, the additional conditions “data stock” (equation 2.10) and “Machine Learning system efficiency” (equation 2.9) need to be defined; these conditions affect variables  $U_{ML}$ ,  $c_{BIAS}$ ,  $r_{ML}$  and, in turn, the price and quantity functions. In order to solve the game by applying backward induction to the game, it is necessary to determine the number of sub-games, which depends on Machine Learning efficiency and how this advances over time.

By assuming that, at each stage of the game, the firm collects non-negative amounts of data from its consumers, it is true that the Machine Learning efficiency will always be higher (or at least equal to) in the next stage of the game (as shown in figure 2.6,  $\eta$  is higher for greater stocks of data). It can be demonstrated from equation 2.9, that the system reaches an efficiency of 100% when the firm collects an infinite amount of data; setting 2.10 equal to infinite, it can be observed that the game played by the firm has an infinite number of stages; the proof of this proposition is reported in the Appendix of this document.




Figure 2.6. Trend of the Machine Learning efficiency versus the stock of collected data

In real applications, however, a Machine Learning system is usually imperfect, meaning that it will always be subject to an accuracy lower than 100%: this is true since it is virtually impossible for the firm to collect unlimited data. Even though this could be possible, however, accuracy could still be affected by the noise contained in the data collected, by some features of the algorithm or by the interaction of both data and algorithm, but could also depend on a strategic decision of the firm that could be satisfied to offer a very high, yet inaccurate, ML system. In this case, the corresponding accuracy can be obtained in a finite sequence of stages that, once again, can be derived from 2.9 and consequently from 2.10. Henceforth, the only case discussed is the one requiring an infinite number of stages and this latter scenario is omitted.

## 2.7 Monopolist firm's behavior

So far, no assumption has been made about the nature of the firm offering the product. For this section, let us assume that it is a monopolist, whose profit function is defined as follows: the revenues depend on the price and the consumers' preferences (which determine demand), while costs are those defined in 2.8 (hence, marginal cost plus the investment in Machine Learning algorithms); moreover, the profit decrease as the firm collects data from its customers. The profit function of the monopolist can thus be written as:

$$\Pi = \begin{cases} (p - c)q - dv - I(\eta_{ALG}) & \text{for stage 1} \\ (p - c)q - dv & \text{for any other stage} \end{cases} \quad (2.11)$$

The profit maximizer firm decides the price and the amount of data that it has to collect from customers. By looking for a subgame perfect equilibrium, the game is solved backward, hence starting with the firm's second problem.

### 2.7.1 Generic stage game of the infinite sequence: optimal values for $p, d$

The profit function can be rewritten using the functions defined in 2.5; the maximization program for the generic  $i$ -th ( $i \neq 1$ ) stage game of the infinite sequence is then:

$$\max_{p,d} \Pi \quad (2.12)$$

The two first-order conditions are computed and later set equal to zero to derive the price and the amount of data collected by the firm. Subsequently, the quantity and the value of the data for the monopolist are calculated by substitution.

$$\begin{cases} p^M = \frac{1}{2}(V + U_{ML} + c - c_{BIAS} + \frac{\beta}{\alpha} \cdot r_{ML}) \\ q^M = \phi[\frac{\alpha}{2}(V + U_{ML} - c - c_{BIAS}) + \beta(\frac{r_{ML}}{2} - v)] \\ v^M = r_{ML} + \beta q^M - \alpha d^M \\ d^M = \frac{1}{\alpha}(\beta q^M + \frac{r_{ML}}{2}) \end{cases} \quad (2.13)$$

The solutions of the subgame are rewritten independently solving the linear system of four equations and four variables above, resulting in the following subgame-perfect Nash equilibria:

$$\begin{cases} p_S^M = \frac{1}{2}(V + U_{ML} + c - c_{BIAS} + \frac{\beta}{\alpha} r_{ML}) \\ q_S^M = \frac{\alpha\phi}{2}(V + U_{ML} - c - c_{BIAS}) \\ v_S^M = \frac{r_{ML}}{2} \\ d_S^M = \frac{1}{2}[\frac{r_{ML}}{\alpha} + \beta\phi(V + U_{ML} - c - c_{BIAS})] \end{cases} \quad (2.14)$$

Three of the four equations depend on the state of the ML system; an exception is the value of the data, which is independent from it. The optimal price for the monopolist has the same form of that fixed by a single-product monopolist firm with linear demand and constant marginal cost, even though in this case, apart from the additional markup arising from the net ML utility, the price is increased by  $\beta/\alpha \cdot r_{ML}$ , since the firm recognizes that the customer values the future expectations about the utility gain due to data sharing. On the other hand, the optimal quantity keeps the same form of the benchmark case, even though multiplied by a different coefficient.

The data are valued accordingly to the gain that the customer expects to earn in the future by sharing her personal data. The amount of shared data depends on the same variables in the formula of the number of goods purchased (even multiplied by *beta* instead of *alpha*); moreover, there is the variable that expresses the value of the data exchanged. It is worth noting that the amount of data exchanged is growing in the performance of the ML system. This behavior is justified by the fact that, given the exponential and asymptotic performance of the system's performance, when the quality is already very high, the system needs many data so that it can improve even a little.

### Towards the dynamic equilibrium

As the Machine Learning system improves its performances, the price and the quantity demanded will shift upwards. The higher quality will also induce the customer to share more data with the platform because of its confidence in the system capabilities. A virtuous circle is triggered, and



the firm will keep collecting data (and therefore improve the quality of its product) until the system reaches an accuracy of 100%.

The convergence speed towards perfect accuracy always depends on the previous state of the system. For sufficiently large time horizons, the monopolist is expected to reach this result with any general level of quality, but of course a better starting point is essential to speed the entire process up.

### 2.7.2 Stage 1: optimal value for the investment

In the first stage of the game, the firm determines the optimal value for the investment; the game is solved by setting the profit function of the monopolist equal to zero, meaning that the optimal investment is equal to the sum of all the discounted profit that the firm earns in every stage different from the first, using the discount factor  $\delta \in [0; 1)$ . Mathematically:

$$I(\eta_{ALG})^M = \sum_{t=1}^{\infty} \left\{ \frac{1}{(1+\rho)^t} \cdot [(p_{St}^M - c)q_{St}^M - d_{St}^M v_{St}^M] \right\} \quad (2.15)$$

From the value of the optimal investment it is also possible to ascertain the optimal efficiency value of the algorithm that the monopolist should implement to maximize its profits; in this sense, it is sufficient to replace the value just obtained in the equation 2.7 and make the calculation:

$$\eta_{ALG} = \frac{1}{\lambda} \cdot \log \left( \frac{A + K - I(\eta_{ALG})}{K} \right) \quad (2.16)$$

### 2.7.3 Consumer surplus

The consumer welfare is measured by the consumer surplus, corresponding to the surface of the area under the demand curve and above the market price. In a market with a monopolist firm, the maximum and the minimum prices that the consumer is willing to pay for the product are known, and respectively equal to  $p_{MAX} = p(q=0) = V + U_{ML} - c_{BIAS} + \beta d^M$  and  $p_{EQ} = p^M = 1/2(V + U_{ML} + c - c_{BIAS} + \beta/\alpha \cdot r_{ML})$ . The definite integral can be computed and it results equal to:

$$\begin{aligned} CS_M &= \int_{p_{EQ}}^{p_{MAX}} q dp \\ &= \frac{1}{2}(V + U_{ML} - c - c_{BIAS})(1 + \beta^2 \phi) \left[ \alpha(V + U_{ML} - c_{BIAS}) + \frac{\beta}{2} r_{ML} \right] \\ &\quad + \frac{\alpha \phi}{2} \left[ \frac{1}{2}(V + U_{ML} - c - c_{BIAS})(1 + \beta^2 \phi) \right]^2 \end{aligned} \quad (2.17)$$

The consumer surplus ranges across different values according to the performances of the Machine Learning algorithm, so that it will differ at each stage of the game different from the first (it will start from the lowest possible value, then reach the maximum value and be constant from there on). If the consumer is uncertain about what is the status of the consumer respect to the ML efficiency, a generic probability distribution is introduced and associated to each of the efficiency states, hence to each of the possible values of the consumer surplus. The maximum value is reached when the additional utility from ML features is maximum ( $U_{ML} = U_{ML}^{MAX}$ ) and when there is no bias ( $c_{BIAS} = 0$ ); conversely, the minimum value is reached when the additional utility from ML features is absent ( $U_{ML} = 0$ ) and when the bias is maximum ( $c_{BIAS} = c_{BIAS}^{MAX}$ ); the two values are therefore those reported below:

$$CS_M^{MAX} = \mathbb{E}[\eta = 1.0] \left\{ \frac{1}{2}(V + U_{ML}^{MAX} - c)(1 + \beta^2\phi) \left[ \alpha(V + U_{ML}^{MAX}) + \frac{\beta}{2}r_{ML} \right] + \frac{\alpha\phi}{2} \left[ \frac{1}{2}(V + U_{ML}^{MAX} - c)(1 + \beta^2\phi) \right]^2 \right\} \quad (2.18)$$

$$CS_M^{MIN} = \mathbb{E}[\eta = \eta_{MIN}] \left\{ \frac{1}{2}(V - c_{BIAS}^{MAX} - c)(1 + \beta^2\phi) \left[ \alpha(V - c_{BIAS}^{MAX}) + \frac{\beta}{2}r_{ML} \right] + \frac{\alpha\phi}{2} \left[ \frac{1}{2}(V - c_{BIAS}^{MAX} - c)(1 + \beta^2\phi) \right]^2 \right\} \quad (2.19)$$

## 2.8 Social optimizer's behavior

Differently from the previous section, it is now assumed that it is not the monopolist firm who offers the product, but rather from a social optimizer firm which has the same characteristics hypothesized for the monopolist firm, except for its goal, that is to balance the interests of both those buying the good and those selling it. In this setting, consumers and the monopolist have the same weight in the measure of total welfare:

$$W = \begin{cases} CS + (p - c)q - dv - I(\eta_{ALG}) & \text{for stage 1} \\ CS + (p - c)q - dv & \text{for any other stage} \end{cases} \quad (2.20)$$

The social optimizer plays the same multi-stage game as the one played by the monopolist in section 2.7, but since its goal is different from just maximizing profit, the resulting quantities are expected to be different; once again, the functions defined in 2.5 are used to determine the optimal solution. The solution proposed below is the first-best one, meaning that the profit constraint ( $\Pi \geq 0$ ) is not necessarily satisfied.

### 2.8.1 Generic stage game of the infinite sequence: optimal values for p, d

Again, the firm's strategic variables are p and d, and the maximization program for the generic  $i$ -th ( $i \neq 1$ ) stage game of the infinite sequence is then:

$$\max_{p,d} W \quad (2.21)$$

Once calculated, the two first-order conditions are set equal to zero in order to determine the optimal values of price and data quantity according to the social optimizer; these values are then substituted in the functions of quantity and data value, resulting in:

$$\begin{cases} p^W = c + \frac{\beta}{\alpha}v^W \\ q^W = \phi[\alpha(V + U_{ML} - c_{BIAS} - p^W) + \beta(r_{ML} - v^W)] \\ v^W = r_{ML} + \beta q^W - \alpha d^W \\ d^W = \frac{\beta}{\alpha}q^W \end{cases} \quad (2.22)$$

The solutions of the subgame are rewritten independently solving the linear system of four equations and four variables above, resulting in:

$$\begin{cases} p_S^W = c + \frac{\beta}{\alpha}r_{ML} \\ q_S^W = \phi[\alpha(V + U_{ML} - c - c_{BIAS}) - \beta r_{ML}] \\ v_S^W = r_{ML} \\ d_S^W = \beta\phi(V + U_{ML} - c - c_{BIAS} - \frac{\beta}{\alpha}r_{ML}) \end{cases} \quad (2.23)$$

The price set by the social optimizer, unlike that set by the monopolist, depends exclusively on two variables: (1) the marginal cost of the product and (2) the personal benefit that the customer expects to earn from ML; this second element, multiplied by the coefficient  $\beta/\alpha$  imply that the price is higher if data have higher complementarity with the data and that it will be lower otherwise. The parallel with the monopoly price allows noticing also how the price that maximizes the collective welfare is independent of the quality of the machine learning system, given that the variables  $U_{ML}$  and  $c_{BIAS}$  are not included. Thereupon, in the case of the monopolist, the price will grow over time (in fact, depending on quality), while in the case of the social optimizer it remains constant. The consumer can, therefore, benefit from a higher quality product by paying the same price at all times.

Regarding the quantity of product requested, unlike the quantities of monopoly, here also the variable  $r_{ML}$ , multiplied by the coefficient  $\beta$ : the higher the complementarity of the data with the product, the more significant the impact (negative) on the quantities requested.

The value assigned to data and their optimal amount depends on the same variables that influenced the monopoly solution, but in this case, they are multiplied by different coefficients. Also in the case of the “data” product, the price does not depend on the quality of the system in which the data are used, but only on the private benefits expected by the customer who decides to share their data.

## 2.8.2 Consumer surplus

If prices and quantities are fixed by the social optimizer, the maximum and the minimum prices that the consumer is willing to pay for the product are known, and respectively equal to  $p_{MAX} = p(q = 0) = V + U_{ML} - c_{BIAS} + \beta d^W$  and  $p_{EQ} = p_S^W = c + \beta/\alpha \cdot r_{ML}$ . The definite integral can be computed and it results equal to:

$$\begin{aligned} CS_W &= \int_{p_{EQ}}^{p_{MAX}} q dp \\ &= (V + U_{ML} - c - c_{BIAS} - \beta/\alpha \cdot r_{ML})(1 + \beta^2 \phi) \alpha \phi (V + U_{ML} - c_{BIAS}) \\ &\quad + \frac{\alpha \phi}{2} \cdot [(V + U_{ML} - c - c_{BIAS} - \beta/\alpha \cdot r_{ML})(1 + \beta^2 \phi)]^2 \end{aligned} \quad (2.24)$$

The considerations about the consumer surplus exposed in section 2.7.3 still hold, including the one about the probability distribution; again, the maximum value is reached when the additional utility from ML features is maximum ( $U_{ML} = U_{ML}^{MAX}$ ) and when there is no bias ( $c_{BIAS} = 0$ ); conversely, the minimum value is reached when the additional utility from ML features is absent ( $U_{ML} = 0$ ) and when the bias is maximum ( $c_{BIAS} = c_{BIAS}^{MAX}$ ); the two values are therefore those reported below:

$$CS_W^{MAX} = \mathbb{E}[\eta = 1.0] \left\{ \left( V + U_{ML}^{MAX} - c - \frac{\beta}{\alpha} \cdot r_{ML} \right) (1 + \beta^2 \phi) \alpha \phi (V + U_{ML}^{MAX}) + \frac{\alpha \phi}{2} \left[ \left( V + U_{ML}^{MAX} - c - \frac{\beta}{\alpha} \cdot r_{ML} \right) (1 + \beta^2 \phi) \right]^2 \right\} \quad (2.25)$$

$$CS_W^{MIN} = \mathbb{E}[\eta = \eta_{MIN}] \left\{ \left( V - c_{BIAS}^{MAX} - c - \frac{\beta}{\alpha} \cdot r_{ML} \right) (1 + \beta^2 \phi) \alpha \phi (V - c_{BIAS}^{MAX}) + \frac{\alpha \phi}{2} \left[ \left( V - c_{BIAS}^{MAX} - c - \frac{\beta}{\alpha} \cdot r_{ML} \right) (1 + \beta^2 \phi) \right]^2 \right\} \quad (2.26)$$

### 2.8.3 Stage 1: optimal value for the investment

The first stage of the game is solved by setting the total welfare function of the social optimizer equal to zero and deriving from it the optimal value for the investment:

$$I(\eta_{ALG})^W = \sum_{t=1}^{\infty} \left\{ \frac{1}{(1+\rho)^t} \cdot [CS_t + (p_{St}^W - c)q_{St}^W - d_{St}^W v_{St}^W] \right\} \quad (2.27)$$

where  $\rho \in [0; 1)$  represents the discount factor chosen by the social optimizer to actualize future profits. The value of the investment for the social optimizer differs from that of the monopolist because the former is also considering the consumer surplus when computing the optimal value. Due to the complexity of the equations computed in the second stage of the game, the result above is left as is and not rewritten with the proper substitutions. Again, from the value of the optimal investment it is possible to ascertain the optimal efficiency value of the algorithm that the social optimizer should implement to maximize the total welfare; in this regard, just replace the value just obtained in 2.7 and make the calculation; the equation will be the same as equation 2.16.

### 2.8.4 Generalization of the total Welfare function

The approach described so far was based on a welfare function that accounts of consumer surplus and profit using the same weights. The social optimizer can, however, be more inclined to protect the interests of consumers rather than profits. The additional analysis proposed in this section is based on Baron and Myerson (1982) and assumes that the welfare function as following:

$$W = CS + k\Pi \quad (2.28)$$

where  $0 \leq k \leq 1$ . The mathematical derivation of the solution is reported in the Appendix; the solution of the subgame result in:

$$\begin{cases} p_k^W = \frac{k}{2k-1}c + \frac{k}{(2k-1)^2}\frac{\beta}{\alpha}r_{ML} + \frac{k-1}{2k-1}(V + U_{ML} - c_{BIAS} + \frac{\beta}{\alpha}r_{ML}) \\ q_k^W = \frac{\phi}{2k-1}[\alpha k(V + U_{ML} - c - c_{BIAS}) + \frac{-2k^2+5k-2}{2k-1}\beta r_{ML}] \\ v_k^W = r_{ML} - \frac{k-1}{2k-1}r_{ML} \\ d_k^W = \frac{k}{2k-1}\beta\phi(V + U_{ML} - c - c_{BIAS}) + \frac{r_{ML}}{\alpha(2k-1)}(k-1-\beta^2(2k^2-5k+2)) \end{cases} \quad (2.29)$$

Although this approach is more formal, it is only reported for completeness. For the monopolist/social optimizer solution comparison it is used the case reported in section 2.8.1, which is also characterized by a more straightforward algebra. The same procedure reported in the same section, however, can also be used for comparing the monopolist's behavior with this social optimizer's behavior.

## 2.9 Solution comparison: Monopolist VS Social Optimizer

The table 2.9, reported below, summarizes the main findings in terms of prices and quantities for both the monopoly solution and the social optimizer solution:

It is then possible to compare the results in order to determine which alternative is better from a consumer perspective. By assumption, the quantities compared are from the same  $i$ -th stage of the game (e.g., they represent the generic solution of the  $i$ -th subgame), whether played by the monopolist or by the social optimizer.

	Monopolist	Social Optimizer
<b>Price</b>	$\frac{1}{2}(V + U_{ML} + c - c_{BIAS} + \frac{\beta}{\alpha}r_{ML})$	$c + \frac{\beta}{\alpha}r_{ML}$
<b>Demand</b>	$\frac{\alpha\phi}{2}(V + U_{ML} - c - c_{BIAS})$	$\phi[\alpha(V + U_{ML} - c - c_{BIAS}) - \beta r_{ML}]$
<b>Data value</b>	$\frac{r_{ML}}{2}$	$r_{ML}$
<b>Data shared</b>	$\frac{1}{2} \left[ \frac{r_{ML}}{\alpha} + \beta\phi(V + U_{ML} - c - c_{BIAS}) \right]$	$\beta\phi(V + U_{ML} - c - c_{BIAS} - \frac{\beta}{\alpha}r_{ML})$

Table 2.1. Monopoly VS Social Optimizer solutions

In terms of product's price, the condition according to which the social optimizer sets lower prices than the monopolist can easily be derived:

$$r_{ML} > -\frac{\alpha}{\beta}(V + U_{ML} - c - c_{BIAS}) \quad (2.30)$$

This condition is always verified when the variable  $r_{ML}$  is non-negative, meaning that the social optimizer is able to offer its product to a lower price.

In the case of the quantity demanded, the comparison of the two quantities leads to the condition:

$$r_{ML} < \frac{V + U_{ML} - c - c_{BIAS}}{2\beta} \quad (2.31)$$

that states that the quantity demanded when the product is offered by the social optimizer is higher than the one demanded in the case of a monopolist when (1) the private benefit arising from data sharing is smaller than the quantity on the left side and at the same time (2), given that  $r_{ML}$  is always a positive quantity, benefits  $(V + U_{ML})$  need to be higher than the costs  $(c + c_{BIAS})$ , and  $\beta$  has to be a positive quantity.

In the case of the data value, the comparison is straightforward: the social optimizer assigns twice the value of that given by the monopolist, and in both cases, the value only depends on  $r_{ML}$ . The social optimizer solution is to be preferred since values more the information that the consumer shares with the platform.

In terms of the amount of data shared, the following equation can be derived from the comparison of the two quantities:

$$r_{ML} > (V + U_{ML} - c - c_{BIAS}) \frac{\alpha\beta\phi}{1 + 2\beta^2\phi} \quad (2.32)$$

meaning that the social optimizer requires fewer data than the monopolist if the personal benefit is higher than the quantity on the right side of the inequality (which is an always-positive quantity). In any other case, the monopoly solution is to be preferred, since it corresponds to a lower rate of data exchange. From the firm perspective, however, the solution that requires fewer customer data will converge to the perfect accuracy in a higher number of periods, meaning that the quantities of the following periods will be affected by this.

## 2.10 A numerical example

This section is dedicated to an application of the model proposed above, in which the numerical values are used for the sole purpose of illustrating the differences that emerge by comparing the monopoly solution with the social optimizer solution. The values are listed below:

$V = 10\$$	$\eta_{DATA} = 0.75$
$c = 5\$$	$\rho = 0.1$
$r_{ML} = 3$	$\lambda = 1$
$\alpha = 1$	$K = 35,000\$$
$\beta = 0.80$	$A = 500\$$

The trend of consumers over time is assumed to evolve according to the logistic equation, or Verhulst model:

$$N(t) = \frac{N_{MAX} N_0 e^{rt}}{N_{MAX} + N_0(e^{rt} - 1)}$$

where

$N_{MAX}$  is the limiting value of  $N(t)$  and assumed equal to 1,000,

$N_0$  is the initial value of  $N(t)$  and assumed equal to 10,

$t$  is the considered period, and

$r$  is the growth rate, assumed equal to 0.2.

It is assumed that the firm has trained the algorithm before being put on the market and that at the instant 0 the following situation occurs:

$$D^T = 1,$$

$$U_{ML}^0 = 10(1 - e^{-1}) = 6.32, \text{ and}$$

$$c_{BIAS}^0 = 10 \cdot e^{-1} = 3.68.$$

### 2.10.1 Monopolist solution

The monopolist solution is computed using the equations derived in section 2.7.

**t = 1**

$$N_1 = \frac{10,000 e^{0.2}}{1,000 + 10(e^{0.2} - 1)} = 12.19$$

$$p^M = \frac{1}{2}(10 + 6.32 + 5 - 3.68 + 0.8 \cdot 3) = 10.02\$$$

$$q^M = \frac{2.78}{2}(10 + 6.32 - 5 - 3.68) = 10.619 \sim 11 \text{ units}$$

$$v^M = \frac{3}{2} = 1.5$$

$$d^M = \frac{1}{2}[3 + 0.8 \cdot 2.78(10 + 6.32 - 5 - 3.68)] = 9.996$$

$$CS = 532.13$$

$$D^1 = 1 + 12.19 \cdot 9.996 \sim 121$$

$$U_{ML}^1 = 10(1 - e^{-0.75 \cdot 121}) = 10$$

$$c_{BIAS}^1 = 10 \cdot e^{-0.75 \cdot 121} = 0$$

**t = 2**

$$\begin{aligned}
 N_2 &= \frac{10,000 e^{0.4}}{1,000 + 10(e^{0.4} - 1)} = 14.85 \\
 p^M &= \frac{1}{2}(10 + 10 + 5 + 0.8 \cdot 3) = 13.7\$ \\
 q^M &= \frac{2.78}{2}(10 + 10 - 5) = 20.85 \sim 21 \text{ units} \\
 v^M &= \frac{3}{2} = 1.5 \\
 d^M &= \frac{1}{2}[3 + 0.8 \cdot 2.78(10 + 10 - 5)] = 18.18 \\
 CS &= 1,322.45 \\
 D^2 &= 1 + 12.19 \cdot 9.996 + 14.85 \cdot 18.18 \sim 391 \\
 U_{ML}^2 &= 10(1 - e^{-0.75 \cdot 391}) = 10 \\
 c_{BIAS}^2 &= 10 \cdot e^{-0.75 \cdot 391} = 0
 \end{aligned}$$

The amount of data collected in the first two stages of the game is sufficient to cancel the bias and to gain the user the maximum additional utility. It is viable to proceed with the calculation of the investment in algorithms that maximize the profit of the monopolist.

$$I(\eta_{ALG}) = \frac{(10.02 - 5) \cdot 11 - 1.5 \cdot 10}{1.1^1} + \frac{(13.7 - 5) \cdot 21 - 1.5 \cdot 18.18}{0.1 \cdot 1.1^2} \sim 1,310\$$$

which leads to the computation of the optimal algorithm efficiency for the monopolist:

$$\eta_{ALG} = -\frac{1}{1} \cdot \log \left( \frac{500 + 35,000 - 1,310}{35,000} \right) = 10.16\%$$

### 2.10.2 Social Optimizer solution

The social optimizer solution is computed using the equations derived in section 2.8.

**t = 1**

$$\begin{aligned}
 N_1 &= \frac{10,000 e^{0.2}}{1,000 + 10(e^{0.2} - 1)} = 12.19 \\
 p^W &= 5 + 0.8 \cdot 3 = 7.4\$ \\
 q^W &= 2.78(1(10 + 6.32 - 5 - 3.68) - 0.8 \cdot 3) \sim 15 \text{ units} \\
 v^W &= 3 \\
 d^W &= 0.8 \cdot 2.78 \left( 10 + 6.32 - 5 - 3.68 - \frac{0.8}{1} \cdot 3 \right) = 11.65 \\
 CS &= 806.84 \\
 D^1 &= 1 + 12.19 \cdot 11.65 \sim 143 \\
 U_{ML}^1 &= 10(1 - e^{-0.75 \cdot 143}) = 10 \\
 c_{BIAS}^1 &= 10 \cdot e^{-0.75 \cdot 143} = 0
 \end{aligned}$$

**t = 2**

$$\begin{aligned}
N_2 &= \frac{10,000 e^{0.4}}{1,000 + 10(e^{0.4} - 1)} = 14.85 \\
p^W &= 5 + 0.8 \cdot 3 = 7.4\$ \\
q^W &= 2.78(1(10 + 10 - 5) - 0.8 \cdot 3) \sim 35 \text{ units} \\
v^W &= 3 \\
d^W &= 0.8 \cdot 2.78 \left( 10 + 10 - 5 - \frac{0.8}{1} \cdot 3 \right) = 21.35 \\
CS &= 3,653.03 \\
D^2 &= 143 + 15 \cdot 21.35 \sim 463.25 \\
U_{ML}^2 &= 10(1 - e^{-0.75 \cdot 463.25}) = 10 \\
c_{BIAS}^2 &= 10 \cdot e^{-0.75 \cdot 463.25} = 0
\end{aligned}$$

The amount of data collected in the first two stages of the game is sufficient to cancel the bias and to gain the user the maximum additional utility. it is possible to proceed with the calculation of the investment in algorithms that maximize the profit of the social optimizer.

$$I(\eta_{ALG}) = \frac{806.84 + (7.4 - 5) \cdot 15 - 3 \cdot 11.65}{1.1^1} + \frac{3,653.03 + (7.4 - 5) \cdot 35 - 3 \cdot 21.35}{0.1 \cdot 1.1^2} \sim 31,090\$$$

which leads to the computation of the optimal algorithm efficiency for the monopolist:

$$\eta_{ALG} = -\frac{1}{1} \cdot \log \left( \frac{500 + 35,000 - 31,090}{35,000} \right) = 89.96\%$$

### 2.10.3 Solutions comparison and discussion

The numerical example shown in this section is well suited to draw some considerations of a more or less generalizable nature. Firstly, it can be noted that: (1) the monopoly price is always higher than the price of welfare maximization, (2) the monopolist's quantity is always less than the amount of welfare maximization, (3) consumer data has less value if the firm is a monopolist, and (4) the amount of data collected by the social optimizer firm is higher than that collected by the monopolist.

As expected, at every stage of the game the consumer surplus in the case of the social optimizer firm is more substantial than what the consumer would get if she bought the product from a monopolist firm.

On a par with other factors, the investment of the social optimizer is far more significant than the one the monopolist would commit to, since the former invests sufficiently to achieve an algorithmic efficiency of almost 90% while the monopolist's accuracy level is about 10%. Both types of firms are still able to extract most of the additional value (then converted into quality) from the data, and both can obtain the maximum net benefit from ML within the second stage of the game.

In conclusion, it should not be excluded that the results obtained are these by virtue of the arbitrariness of the values to be assigned to the variables. These values are associated with the maximum efficiency of the ML algorithm already starting from the second stage of the game, and therefore make it impossible to compare the trend of the monopoly solution and the welfare maximization solution in a higher number of periods.



## 2.11 Chapter conclusions

The second chapter was dedicated to a microeconomic model of customer-firm interaction, with the firm offering a smart product (e.g., powered by ML algorithms capable of enhancing the product) and customers required to share their data to be able to use the product. The interplay of the two agents has been modeled as taking place in two interconnected markets, one dedicated to the good and one dedicated to the data. This setting makes both agents customers and sellers at the same time.

Since the AI improves as it collects more data from its customers, the interplay has been discussed as a dynamic game with an infinite number of stages. The problem can, however, be reduced to a game with a finite number of stages by constraining the ML performance efficiency to a value smaller than 1.0, this decision corresponds, in fact, to a more realistic scenario.

Data are an essential ingredient for the firm to offer higher quality to the customer since they are used to feed an ML algorithm that in the end allows users to gain additional utility from the product consumption. To some extent, in this case, the economics of ML was seen as the economics of data. ML is still necessary to enhance the expected value of data, meaning that both customer and firm can internalize a share of the value created by the algorithm's output.

Assuming the firm to be first a monopolist and later a social optimizer allows to evaluate which would be, from a consumer perspective, the solution that maximizes her welfare. It has been shown that, for the social optimizer to be the best solution, some conditions on the variables need to be respected. Even though the model has not been tested empirically, it can still be used to describe (making proper simplifications) real-world cases.

An interesting finding is that the value assigned to the data shared only depends by the expectations of the customer about the future utility arising from its decision to share personal data; in this case, the social optimizer is a better solution since it allows customers to value data the double, compared to the monopolist firm. The social optimizer also sets its price independently from the current performance level of the ML system, making this value constant over time (this is not true for the monopolist, who is expected to increase the price of its product as it improves over time).

This model is meant to be a starting point for a further discussion about the role of the consumer when interacting with AI-powered products. In the next chapter, some of the variables and effects presented in this chapter are furtherly broken down and discussed, in order to provide useful insights to the policymaker interested in dealing with issues like data privacy and consumer manipulation regarding biased products and loss of decisional autonomy.

## Chapter 3

# Case study discussions and normative recommendations for policymakers

The findings from the previous chapter show how the customer-firm interaction on two interconnected markets has consequences on how both parties decide how to play with the other one. The understanding of customers changed dramatically thanks to the firm's capability to collect (big) data, and this could result in a change of the power balance between customers and firm under many aspects.

The remainder of this chapter is arranged as follows; it starts with the discussion of two case-studies, by introducing of two companies famous for their real market-applications of Machine Learning in two different industries: Netflix, the popular video streaming service and Amazon, the even-more-popular online shopping retailer. It is later questioned if, and under which assumptions, the proposed model fits the way the two firms conduct their business. The second segment of the chapter describes some issues arising from the implementation of AI-features of products which are relevant for policymakers, to safeguard the interests of consumers. They are issues related to customer data privacy and customer manipulation (whether voluntary or not). Spark this discussion is vital because, by designing their regulatory environment as well as directing public expenditure, countries can accelerate the development of AI and provide the country with a comparative advantage on this field.

### 3.1 Case study of two digital platforms

Netflix and Amazon are just two examples of an emerging spate of innovative, data-driven followers (and newcomers) who threaten to raise the bar of customer-firm interaction even higher. Consider Google, Airbnb, Uber, Spotify, each of which has achieved billion-dollar valuations in just a few years, or many of the companies introduced in section 1.2.

All of them share a similar data-centric culture that proposes to improve people's lives leveraging on what they are willing to share with the formers.

### 3.1.1 Digital video streaming: Netflix Inc.

Netflix Inc. is an American enterprise that (mainly) offers a Digital Video Streaming (DVS) service in more than 190 countries to more than 137 million subscribers worldwide <sup>1</sup> (Richter, 2018). Its business model is based on a two-sided platform which sees the content producers on one side and customers on the other.

On the users' side, there are direct network externalities, since the content recommendation algorithm implemented by Netflix bases its output on the contents watched by other users, along with the user's chronology. There is a positive feedback loop that works like this (Schepp and Wambach, 2015): as the number of users increases, the amount of data generated increases, allowing the algorithm to make more accurate recommendations. In turn, greater accuracy in the recommendation attracts new users and so on.

The critical aspect of Netflix's business model is the proprietary content recommendation system. To offer the best possible experience, Netflix reconciles in the right proportion of personalization and the proposal of popular titles. Through this strategy, Netflix tries to replicate, to a certain extent, the experience of walking between the shelves of a video library, in which the available contents are continually changing.

The company reportedly tracks what a user streamed, searched for, rated, as well as the time, date, and device. User interactions like browsing or scrolling behavior are recorded as well (Solow, 1957). Such ML algorithms, to elaborate a recommendation, are not limited to considering the contents consumed in the past by the spectator himself, but they implement pattern crossing functionality among similar viewers to offer much more personalized recommendations (Rayna and Striukova, 2016). In doing so, the value perceived by the user for this service is much higher, since she is likely to find a content immediately that she is going to enjoy, without spending too much time in researching it.

Thanks to its first-mover advantage in the DVS industry, Netflix has the most extensive collection of video ratings in the world (Shih et al., 2007), and therefore can customize the user experience better than its competitors. The recommendation allows users to be shown content that will be appreciated, but which would hardly have been discovered by the user himself; this makes the contents that make up the long-tail of the product stock more profitable.

Netflix "learns to order" from its users to offer increasingly accurate recommendations: show the same contents in a different order (or possibly hide some to show them in the future), gives the impression to the user that the catalog is continually evolving, limiting at the same time investments in new content. To contain the costs deriving from the acquisition of video content, Netflix thus stimulates the demand for older and lesser-known content, possibly already present in the catalog.

The company also employs all of its collected data to create plots for original video contents: using consumers' habits, the company is capable of engineer a show that has the right elements to become a phenomenon; in fact, the success rates for Netflix's original shows are much higher success rates of traditional TV shows offered on the platform. The company has been an innovator in this sense, and this value creation process is currently being explored by many other players of the Media and Entertainment industry (The Economist, 2018).

The profiling of the user constitutes, on the one hand, a source of market power for the company and, on the customer side, a not inconsiderable switching cost (Schaefer et al., 2018b), which therefore is encouraged to remain on the platform and not to turn to a new video-on-demand provider. Considering the non-rival, yet excludable nature of the data, economies of scale linked to the information gathered that limit the threat that new entrants could represent can be found as well.

---

<sup>1</sup>As of 18th October 2018

### 3.1.2 Online retailing: Amazon Inc.

Amazon Inc. is the second-biggest public corporation for market capitalization <sup>2</sup> and also among the broadest e-commerce platforms, a big data behemoth that is currently offering its services in a variety of industries, ranging from logistics to finance. Even though these businesses may seem utterly unrelated at first sight, a close observation makes it clear that they all share a common feature: they rely upon a profound knowledge of their customers. The biggest accountable for the company's profitability is the Amazon Web Services (AWS) business unit, a pay-as-you-go cloud service; ironically, AWS is also implemented by Netflix to offer its video streaming service.

The company can rely upon a customer base of more than 300 million active customers worldwide (Duprey, 2018), from which it can be supplied of information such as (Amazon, Inc., 2018): (1) name, address and phone numbers; (2) payment information; (3) delivery details of people to whom purchases have been dispatched or people listed in 1-Click settings; (4) content of reviews and e-mails sent to the company (5) personal description and photograph of the personal profile; (6) voice recordings when the customers interact with Alexa, the company's smart personal assistant. Using ML and data analytics, Amazon can easily drill down into the consumers' history to increase the engagement rate with the platform thanks to targeted recommendations and tailored contents; these tools are used to encourage customers to buy on impulse, hence to spend more. A few popular applications include:

**Personalized Recommendation System** Using a comprehensive, collaborative filtering engine, Amazon analyzes the data listed above to recommend additional products that other customers purchased when buying those same items. Adding a gaming console to the virtual shopping cart, video games for that console purchased by other customers are also recommended for the customer to purchase.

**Book Recommendations from Kindle Highlighting** After acquiring Goodreads in 2013, the company integrated social networking features into some Kindle functions. As a result, Kindle readers can highlight words and notes and share them with others as a means of discussing the book. Amazon regularly reviews words highlighted in Kindle devices to determine what users are interested in learning about; this knowledge is later translated into purchase recommendations.

**Anticipatory Shipping Model** Amazon's patented anticipatory shipping model uses big data for predicting the products you are likely to purchase, when you may buy them and where you might need them. The items are sent to a local distribution center, to be ready for shipping once the customer orders them.

**Supply Chain Optimization** Amazon wants to fulfill orders quickly, so the company links with manufacturers and tracks their inventory. Big data systems choose the warehouse closest to the vendor and the customer, to reduce shipping costs.

### 3.1.3 Cases discussion

The takeaway of the two companies discussion above is that both Netflix and Amazon heavily rely on the knowledge about their customers to enhance and personalize the platform experience so that higher engagement will turn into higher profits. Both companies know their single customers better as long as they keep interacting with them, and they keep to paint a more detailed picture of their purchase patterns, their tastes, and preferences. To some extent, these companies are also capable of inferring customers' behavior outside the platform.

Before proceeding with the analysis, it is essential to verify if the assumptions made for the model hold for the two companies described. Even if Netflix faces the competition of many

---

<sup>2</sup>As of 3rd November 2018

competitors, including Amazon (through the Prime Video streaming platform), Hulu and HBO (both USA-based companies), the firm can be seen (to some extent) as a monopolist so that it can be evaluated if the model presented in this thesis applies to it. The same logic applies to Amazon, which is indeed not the only e-commerce platform existent but is undoubtedly one of the biggest (if not the biggest) and one of the few with such a high global pervasiveness. Both implement ML algorithms that enhances over time due to the data that the customer mandatorily shares with them and charge a price for the services they offer.

The complementarity of both services with data (the  $\beta$  of the model presented in the previous chapter) is expected to be high, just like the added benefit that the customer gains using these platforms rather than two others. Of course, in both cases, bias is relevant as well. In the case of Netflix, especially during the first months of subscription, the platform is likely to recommend contents that are very (if not too) similar to those previously consumed, meaning that there is no space for a variegated recommendation that would be better appreciated by customers. In some extreme cases, the user is induced to give up the benefits of the recommendation to pursue first-person research of the content to consume. The e-commerce platform of Amazon has, on the other hand, a similar problem to this one: the recommendation algorithm is biased in the sense that once the customer purchases (or conducts a research about) a product, she is very likely to be recommended products belonging to the same category of the purchased one, rather than complementary products. Let us assume that a customer buys a TV screen on the platform; after the purchase, the AI is likely to recommend her to buy another TV among, for example, home theater sound systems or Blu-Ray disk readers. To conclude, in both cases the bias is detrimental for the customer experience: since the customer realizes that her decision to share data with the platform is not yielding the expected return (even though the platform can profitably use those data to improve the overall system). This event, of course, will also impact the revenues earned by the platform, since worse recommendations are less likely to translate into purchases.

Concerning pricing strategy, the two cases are considered separately. Netflix charges its users with a flat monthly fee that allows customers to consume as many video contents as they like. According to the model, the firm should charge a progressively-increasing monthly fee, that should settle after the ML system reaches perfect accuracy. In real life, however, the price charged by the company is constant through time, except for some occasional increases that, most of the time, are justified on the grounds of expansion. For this reason, if the model proposed is correct and accurate, the firm is deliberately charging a price that is not optimal for the firm, and that could be

- lower than the average of all prices possibly applicable,
- equal to the average of all prices possibly applicable, or
- higher than the average of all prices possibly applicable.

If the first case reveals to be the correct one, it means that the firm is giving up some of its profits and that the remaining value remains in the hands of the users in the form of consumer surplus. In the second case, it means that the consumer, in the long run, is paying a fair price (a bit higher before reaching the average, a bit lower than it should later). Instead, concerns arise if the applied case is demonstrated to be the third one: it would mean that the customer paid a price that let her with less consumer surplus.

About Amazon, the same logic could either be applied to the single purchase or to the Amazon Prime subscription fee, which gives users access to streaming video, free shipping, and other specific services or discounts. Even though in this case, the environmental complexity is higher, due to the firm offering, it is unlikely to find that a user can rebuy the same item at a higher price on the next day (Amazon processes tens, or even hundreds of sales every second, meaning that in 24 hours the amount of collected data is actually translated into better ML predictions, hence in higher pricing power for the firm). The same Amazon Prime service has a constant price over time, except a few cases when permanent surcharges are made in the name of expansion.

Lastly, the variable  $U_{ML}$  can be interpreted as the variable that incorporates the recommendation process. The usefulness deriving from the recommendations will depend on the accuracy, the variety, the novelty of the recommended items: increasing one of these two factors, increases the utility perceived by the consumer (Ziegler and Lausen, 2009).

## 3.2 Issues arising from smart products and relevant for the regulator

The digression of the first section of this chapter, although qualitative, allows the reader to contextualize the proposed model in real applications, with the goal of fostering a discussion on the implications that ML-powered products can have on consumers. Where not all simplifying model assumptions are met, many outstanding problems and conflicts arise.

Aside from many positive outcomes, developments in ML and AI can generate tensions among firms, consumers and policymakers. Also, due to the variety of participants involved in big data and to the potentially enormous economic profit enabled by ML technologies, the potential associated problems are relevant. Moreover, there is a perceived lack of fairness, especially given the lack of standardization of privacy practices across jurisdictional boundaries. Conscious of the fact that there is nothing like a blueprint for the success of ML that goes to the full advantage of the consumer, the author uses the model proposed as the starting point for discussing at least three issues relevant for the regulator:

- inaccurate assessments and data discrimination (e.g., the presence of bias in the system),
- consumer myopia and her potential loss of decisional power, and
- data privacy, ownership and management.

Treating these issues requires to draw knowledge from very different fields of research, such as information systems theory, consumer psychology, economics, and so on. This variety justifies why these topics are broken down and faced separately in the following sections of this chapter.

## 3.3 Issues related to bias in smart products

AI is a powerful tool to enhance the user's experience and customer loyalty (Yoon et al., 2013), but it must be used cautiously and in the right amount. In the model proposed, bias perturbs in a negative way the quality perceived by the consumer; according to the type of product in which ML is implemented, consumers are willing to have different levels of bias tolerance. If information asymmetry is detrimental to the efficiency of economic activity, providing incorrect information can lead to even worse consequences.

Practitioners in sensitive areas like doctors, judges, accountants could wrongly take information from an AI system as they would do with a trusted colleague; an unconditioned trust could lead to serious implications. To illustrate the phenomena, consider the panel (b) of figure 3.1 there is a dog that could be potentially misidentified as a wolf by an AI algorithm because the AI associated the wolf status to the snow that was girding the animal. The mistake is attributable to the bias present in the data set that was fed to the algorithm (e.g., most of the pictures of wolves were in snow). The worrying thing about this is that (and this is perhaps the biggest problem with AI algorithms, deep learning, machine learning) in the case of black box algorithms, developers who worked on them do not have a clear idea of why a particular output is given to users.

Aside from the research implications of this phenomena, the real world outgrowths of bias are the most important thing. A criminal sentence algorithm, for example, could wrongly mistake an

innocent person (a metaphorical dog) for a felon (a metaphorical wolf), all because of a biased AI. An algorithm used to determine if a person can be granted with a loan could mistake a person with a safe credit profile (a metaphorical dog) for someone with an adverse credit profile (a metaphorical wolf).



Figure 3.1. Example of biased AI: a dog (panel b) misidentified for a wolf (panel a) because of the similarity of the background

### 3.3.1 Nature of the problem

The bias may be present in the system due to three reasons:

- Poor AI training, done with few or insufficiently varied data,
- Introduction of bad-quality data in the system by third-party malicious agents, or
- Use of bias as a strategic variable by the firm.

In the first case, once detected, bias can be artificially wiped out feeding the algorithms with more diverse data. If not directly detected by users or developers, it is reasonable to think that it tends to zero when the firm has collected an infinite amount of data (still, assuming that these data have a proper degree of variety).

The second case is more controversial than the previous. A system left to learn by anyone interacting with it can either be influenced by malicious customers or by competitors. Done, of course, to the detriment of consumers who use the service correctly and expect additional utility from the intelligent system.

In the first case, users may have the goal to notch the system performance, rather than making them improve. For instance, research shows a causal impact of online user-generated information on real-world economic outcomes: additional content on Wikipedia pages about Spanish cities increases the number of nights spent in these cities (Hinnosaar et al., 2017). Hence, the study proves the positive effects of digital public goods to inform customers and affect their choices. It raises concerns on how this could be misused to harm customers. The same Wikipedia has been proved to be not wholly free from biased information in a particular context (Greenstein and Zhu, 2012).

In the second case, competitors aiming at incentivizing customers to purchase substitute products could act similarly. This kind of attack goes under the name of shilling attack, which attempts to manipulate the system's recommendations for a specific item by submitting misrepresented opinions to the system (Lam et al., 2006). Different attacks will have different outcomes accordingly



to the robustness of the ML algorithms; distributed recommenders give a potential solution to this problem. The case of Microsoft’s AI, Tay, is emblematic. In 2016, it took less than 24 hours for Twitter users to corrupt an innocent AI chatbot, that was meant to get smarter the more the people chatted with it. As soon as Twitter users started tweeting any racist, misogynistic or any other inappropriate phrase, the situation escalated (the AI started conversations using very inappropriate contents) and the experiment shut down by Microsoft (Perez, 2016). The adage “garbage in, garbage out” has never been so accurate.

The third scenario is a direct consequence of the fact that ML used for recommendations purposes serves the interests of consumers as well as that of its provider. If the recommendation is a strategic variable (and therefore not exogenous, as in the former two scenarios), the firm may have an incentive to alter its recommendations deviate (hence, violate customers’ trust), hence to steer customers and increase its profits. This conduct is particularly true when customers do not internalize the differences in platform costs (e.g., subscription-based platforms) (Bourreau and Gaudin, 2018). Multi-sided platforms may also have an incentive to distort their outputs (hence, bias them) towards their preferred output because of spillovers resulting from this action, like revenues from advertising markets (Burguet et al., 2015) or the creation of new equilibria that benefit the users (Casadesus-Masanell and Halaburda, 2014), even when ML is not directly involved.

The goal of the purposefully-biased firm would consequently be attracting more customers or offer alternate products. In the second case, for instance, the firm could either decide to offer a very-high expensive product that is likely to be enjoyed by customers or offer a cheaper product and induce customers to believe (through recommendations) that they are going to enjoy it. The firm could also convey signals that reveal some relevant and meaningful information to consumers, thus reducing uncertainty and facilitating a purchase or an exchange (Connelly et al., 2011), even though the consumer did not plan this in the first place. Firms may leverage the consumers’ trust and attitude to induce impulse buying behavior. An example can help to clarify the concept.

Web mapping services like Google Maps are used by millions of people worldwide for road navigation, both with vehicles and on foot; in addition to this function, the software is used to report the commercial activities in the area in which the individual is located. As the company’s revenues come from the sale of advertising space present in the various services offered, the platform could accept to charge a fee to the company that needs to be advertised, to ensure that users who find themselves using the service in that area will be steered in order to pass close to the advertised party itself. This fee would hypothetically be higher than the additional profit the firm could realize providing a perfect ML-powered service. In this case, the bias takes the form of the customer “hijacking”: the user is, in any case, able to move from point A to point B as she had set out to do, but instead of choosing the optimal route, she is liable to pay a small price for an alternative route (figure 3.2).




Figure 3.2. Depiction of a biased route (red line) suggested by a web mapping service instead of the optimal one (green line).

The same logic is similarly applicable to video or music streaming services. An emerging artist, for example, could sustain higher advertising costs than the “standard” one to make her product be



shown to those who, as a result of Machine Learning, are labeled as not interested in that content in particular. The user affected by this phenomenon will also see among the songs/videos suggestions that deviate from her preferences, thus perceiving a slightly lower experience of use. Empirical evidence from an e-commerce context shows that, when consumers receive personalized biased recommendations, they are more inclined to choose sub-optimal products, and paradoxically they perceive greater confidence in their choice, less extensive product research and a lower perceived cognitive effort (Xiao and Benbasat, 2018).

The consumer must be aware of the fact that the algorithms to which they are entrusting may give a distorted output compared to the optimal one. Educating the customer precludes them from engaging in behaviors that may lead to undesirable outcomes. It is compelling for regulators and policymakers to instruct consumers about risks of manipulation. About industry leaders, their voluntary adoption of warning tools on their website could make consumers look at them as helpful and honest, instilling trust in the consumer and potentially enhancing customer engagement (Xiao and Benbasat, 2015). Conversely, overstepping the boundary of what buyers acknowledge fair use may unleash a backlash with significant implications for the firm.

### 3.3.2 Recommendations for the policymaker

The digression of the previous sections has a simple bottom-line: AI is a powerful tool, and in many cases is a model for efficiency. People who are not entirely aware of how these systems work could be led to trust smart machines fully, but this does not mean that they could not be wrong. Here, the role for the policymaker is double: regardless of the factors causing the bias (built-in, injected or developed as an unintended consequence), it is its responsibility to foster knowledge about the risks of bias among all those involved in AI-supported decision making and provide a robust legal framework for algorithmic auditing.

Standards for accountability, transparency, the possibility of legal appeal against AI systems are an excellent starting point. Additionally, periodical reviews, a data filtering system before uploading to the data pool; even liability for the firm in the case of incorrect behavior contribute to grant fairness to the final users.

In this sense, the European Union (EU) is moving towards the right direction; the Article 22 of the General Data Protection Regulation (GDPR) gives EU citizens the right to question and oppose to “decisions that affect them that have been made on a purely algorithmic basis”. However, the same cannot be said for many other Countries worldwide: not even the U.S.A. can rely on overarching legislation about data security or consumer privacy (Jin, 2018), even though their data protection laws are considered adequate by the EU. It is clear that even other governments should address the problem and follow the footsteps of the EU.

Explainability is particularly crucial for black box algorithms. The regulator (equipped with the proper independence and the right technical skills) should create its pool of AI experts that should be able to open-up and reverse-engineer the models used by companies or public organizations if its functioning stirs concerns. If the EU, through the GDPR, is already requiring companies to create “explanations” for their models’ internal logic, the DARPA in the USA is developing the Explainable AI program, aimed at interpreting the deep learning that powers drones and intelligence-mining operations (Gunning, 2018).

For the most delicate situations, it would be in consumers’ interest to “cap the bias” by law, avoiding that firms with poorly efficient solutions make their way to the customer. When the AI is dealing with a process with very little tolerance for failure, a human should oversee the intelligent machine and adopt corrective actions when it mistakes. Over time, after the AI learned from its mistakes, it will make the human correction redundant, allowing the firm to put the smart algorithms at work on their own.

### 3.4 Impact of AI recommendations over consumers’ decisional power

By scaling, ML-powered products lower the cost of predictions, which then become more accessible to obtain and more abundant in volume. Being exposed to a higher rate of predictions causes a consumer to apply the decision-making on aspects that only previously accepted the default option. Hence, ML allows consumers to perceive higher utility better matching demand and supply in many ways, including personalization of contents and time savings. For instance, e-commerce platforms are associated with less concentrated sales distributions if compared to traditional channels (Brynjolfsson et al., 2011). Empirical evidence shows that this is also true for music streaming platforms, for which it has also been reported a higher long-run rate of consumption if compared to brick and mortar business models (Datta et al., 2017). Nevertheless, the risk to become a customer that makes her choices using the autopilot can be detrimental for the consumer power, so that it is crucial to raise awareness about how to balance the weight of machine outputs with the weight of human decisions. The point is that usually machines do not see the bigger picture and base their behavior upon incomplete information. While Google Maps provides the shortest route to a destination, this output does not incorporate information like the necessity of fuel for the vehicle or the necessity of rest for the user. Humans, in the other hand, possess their knowledge about why they are doing something, and this gives them the personal touch (a machine can hardly give that) and, of course, the ability to override the output provided by the machine.

Here, the author implies two things: (1) the consumer will always see only part of the bigger picture, because some content will be hidden from the ML system (in the case of video streaming services, the catalog offered at a precise moment will never include all the contents platform, just like an e-commerce platform will tend to hide products that might not please the user) and (2) the consumer, aware of the fact that the ML system will always advise him what she likes, will be led to trust the algorithms’ suggestions, implicitly losing part of its decision-making power. Given the presence of some form of inertia in the purchasing behavior of the consumers, this scenario might be detrimental to consumer welfare.

These implications can be found in many real-world situations where the primary goal in repetitive and relatively unimportant decisions is not to make an optimal choice, but rather to make a satisfying choice that minimizes cognitive effort. This statement is true when: (1) decisions do not involve a degree of risk that does not justify significant decision making effort; (2) consumers made these decisions several times in the past (Hoyer, 1984). To illustrate, if an individual wants to listen to music, watch a film or reach a particular place, she wants to make as little effort as possible, and will, consequently, be more inclined to give away data for profiling and to delegate decision-making autonomy to the algorithms.

What said seems to be confirmed by some empirical data. According to Youtube’s Product Chief, “for 70 percent of the time you watch, you’re riding a chain of recommendations driven by artificial intelligence” (Solsman, 2018). Moreover, Spotify’s playlist Discover Weekly is an AI-powered weekly playlist <sup>3</sup> that generated almost 5bn streams of track since launch (Statista, 2018); in 2016, more than 40% of Spotify’s active users were streaming this playlist (Musically, 2016). Around 23% of consumers interviewed in the US in March 2018 believes that curated playlists are “Very Important” in music streaming services; 28% of them believes that curated playlists are “Somewhat Important” (Statista, 2018). Lastly, back in 2013 Netflix estimated that recommendations drive almost 75% of streaming activity (Vanderbilt, 2013).

Many alternative solutions surround customers for almost each one of their needs. If at first sight, this seems to be good for them, it could become problematic in some situation. For instance, when consumers choose between desirable options, even though they think that put more effort

---

<sup>3</sup>The AI builds a taste profile for each user, based on their past listening history and on similar songs that said user has not yet listened to

in the decision process yields more satisfying outcomes, they experience post-choice discomfort as soon as they have chosen one alternative over others. Individuals experience this feeling because they become attached to their choice options (Carmon et al., 2003). Decisional autonomy leads customers to bear costs pertinent to conflicts, ease of choice, option attachment, choice overload and guilt from choices (André et al., 2018). To quantify the gain (or losses) involved in the decision process, thinking costs must be quantified (Shugan, 1980), as well as emotional or temporal costs (Botti and Hsee, 2010).

Consider a freshly postgraduate student that has to choose between pursuing an academic career or a career in a big company. She identified two options, but she cannot decide which one likes better. So she deliberately takes a long time to choose which options suits her tastes better; as soon as she decides what to do, rather than feeling relieved about having put the conflict to an end, she feels uneasy about her decision, and she is struck by a sense that the other option was more appealing than before choosing.

In this or many other situations, the individual using a smart agent to help him decide may not experience any discomfort if the AI is powerful enough to guide the subject to the right decision. The inquiry to be investigated is what is the price to pay for this relief. The answer requires to quantify the degree of autonomy that the customer had to give away to the machine.

### 3.4.1 The choice problem and choice overload

Economists agree that the higher the number of options in the choice set of a consumer, the higher the likelihood that she can find a close match to her purchase goal. Paradoxically, too much variety can be detrimental to choice, since this variety corresponds to an increase in the cognitive costs associated with choosing from a vast assortment. In specific circumstances, the costs associated with the time spent seeking for the best option may even be higher than the benefits that option provides, resulting in suboptimal outcomes and unpleasant feelings: travelers, for example, may end up paying a higher price for the same plane ticket (Botti and Hsee, 2010).

Choice overload is, hence, a kind of disutility; even though some meta-analysis proved that more choice is better (Scheibehenne et al., 2010), it cannot entirely be ruled out that choice overload may happen when some preconditions occur. In the proposed model, the variable  $U_{ML}$  incorporates the “choice unload factor” that the consumer can benefit. Of course, the benefit perceived by the client will vary according to (Chernev et al., 2015): (1) the complexity of the decision-making task (time constraints, authorization decision, and so on); (2) the complexity of the set among which the consumer must choose (presence of insignificant, complementary options); (3) degree of uncertainty about a preference; (4) the decision-making objective, that is, the degree by which individuals want to minimize the cognitive effort associated with choice. It follows that, to some extent,  $U_{ML}$  incorporates the amount of power that the consumer delegates to the machine compared to what it holds for itself (variable  $V$ ): other factors being equal, a high  $U_{ML}$  value implies a higher incidence (higher power) delegated at the machine, while a lower  $U_{ML}$  implies less power in the hands of the AI.

Policymakers and researchers generally assume that lowering costs associated with search and decision making is empowering for customers and increases their welfare (Stigler, 1961). Algorithms that can predict a person’s most preferred products or services by passively learning her tastes are quite likely to be found practical and convenient in most settings. However, this requires tradeoffs for customers, since AI technologies may undermine the sense of autonomy and free-will. Consumers are not only guided by hedonism but also by a wish for autonomy (André et al., 2018); this theory can be referred as someone’s ability to “be [one’s] own person, and to be directed by considerations, desires, conditions, and characteristics that are not simply externally imposed upon one, but are part of what can somehow be considered one’s authentic self” (Christman, 2008).

The power of personalized recommendations to influence consumer decision making is explained by so-called perceived personalization, defined as the extent to which the consumer believes that the recommendation system understands and represents personal preferences (Komiak

and Benbasat, 2006). Perceived personalization increases the intention of the consumer significantly to adopt a system of recommendations increasing their cognitive and emotional confidence in these systems. An autonomous car manufacturer would want to avoid causing perceptions among customers that they give up their autonomy by being carried in such a vehicle. This goal could incorporate insurances that consumers may still take control of the car if they want, for instance, to customize traits of the self-driving algorithm (driving style, choice of roads, and so on). Contrarily, if a feeling of struggle and conflict is the key to generating a sense of agency, then the company may paradoxically be better off stressing moral aspects of renouncing to drive a car; for instance, by choosing to let an AI drive the vehicle, the consumer commits to making the roads safer and transportation more energy efficient (André et al., 2018).

In many occasions, academic research observed that recommending agents can potentially steer consumers in a particular direction (Adomavicius et al., 2013; Cosley et al., 2003; Häubl and Murray, 2006), which is not necessarily the optimal one for the customer. At this point, the discussion moves towards the topic of manipulation.

### 3.4.2 The problem of manipulation

The risk associated with smart products is that they could end up by manipulating the consumer (either voluntarily or by choice), becoming too present and invasive in the decision-making process, or even making decisions on behalf of the consumer. A consumer will be more inclined to accept one of the proposals made by the ML system both because she puts trust in its faculty, and because to look for the content she wants precisely at that moment can turn out to be tedious and expensive.

Psychological manipulation can be defined as “any intentional act that successfully influences a person to believe or behavior by causing changes in mental processes other than those involved in understanding” (Faden and Beauchamp, 1986). Manipulation is also defined as “directly influencing someone’s beliefs, desires, or emotions, such that she falls short of ideas for belief, desire, or emotion in ways typically not in her self-interest or likely not in her self-interest in the present context” (Barnhill, 2014).

An agent is said to be manipulative if it does not sufficiently engage or appeal to people’s capacity for reflective and deliberative choice. Two problems arise when dealing with manipulation: (1) it fails to respect people’s autonomy and is an affront to their dignity (by making them instruments of another’s will); (2) if people’s choices are products of manipulation, people might promote the welfare of the manipulator in spite of their own welfare (Sunstein, 2015).

Since manipulation comes in many forms, degrees, and pervasiveness, it is challenging to be regulated by governments. It cannot even be excluded that a benign, knowledgeable manipulator could make people’s lives go better and possibly much better. Paradoxically, people might benefit from being manipulated; people have always relied on someone with a more in-depth knowledge of their own in case they need it. If a person decides to change her diet, she decides to turn to a dietician, whose job is to establish the most suitable diet for the user. Also, in this case, the user loses decision-making power (deciding what to eat is not him anymore, but a third party), but this loss is not considered harmful because the consumer is aware of the fact that the decisions taken by the dietician can result better than those made personally.

However, under realistic assumptions manipulators are unlikely to be either benign or knowledgeable (Sunstein, 2015), making preemptive actions necessary. In the case of AI-powered products (as in the case of the dietician), the consumer has a feedback tool to determine her degree of satisfaction concerning the predictions made by the machine. It can send signals to the machine, which can then carry out corrective actions for the future. Based on the current behavior of the algorithms, the consumer can then decide whether to continue relying on the AI or giving it up.

The model does not incorporate how algorithms may manipulate the consumer, but the variable  $r_{ML}$  indirectly reflects the consequences of this event. If the consumer feels like she is going to be manipulated by algorithms, she will have lower expectations about her future gains, so  $r_{ML}$  will be lower in value.

### 3.4.3 Recommendations for the policymaker

Here, the role of the policymaker is to supervise the work of the company to assess whether the consumer, despite receiving recommendations of quality and in line with her interests, is indeed subject to loss of decision-making power. Again, the policymaker's task is anything but simple to implement, yet essential for guaranteeing consumer protection and proper business-to-business cooperation.

It is necessary to furnish consumers with tools that allow them to understand the *whys* of a recommendation, to build trust in AI-powered products and facilitate demand and supply match. In this way, the customer will be empowered to understand whether she is subject to some form of manipulation by the firm.

The manipulation scenario could be modeled and quantified with a dynamic two-players game, in which the company chooses to make a particular recommendation to the user in each period (this recommendation turns into profit based on how good the recommendation is). Each time the consumer plays, she has to choose between repeating the best move she found so far, that is either cooperating (e.g., accept the recommendation) or not (according to the quality of the recommendations given by algorithms) or trying other moves, which gather information that may lead to even better payoffs (e.g., collect feedback from users so that future recommendations will undoubtedly be more accurate). By analyzing how the consumer is playing versus what would be her optimal solution, it could be possible to determine whether there is some form of manipulation or not.

## 3.5 Data Privacy

When data storage was expensive, and the size of storage limited, data was collected more selectively and decisions made before the collection process. As storage capacities expanded and simultaneously became less expensive, more massive datasets could be collected and stored, allowing for more options and flexibility in analysis.

In 2012, Target, a retailer of grocery and home goods in the U.S.A., sent coupons for baby clothes and cribs to a teenager before her family knew she was pregnant. The predictive analysis that resulted in the offer mailing was based on the shopping habits of those enlisted in Target's baby registry. Leveraging their purchase and search patterns, analysts at Target created a list of 25 products that could indicate a woman was pregnant (whether or not enrolled in the baby registry), such as special lotions or pre-natal vitamins. This incident is a notorious example of how invasive data analysis has become ([Duhigg, 2012](#)).

In the information age, privacy is one of the most interesting problems. People leave a trail of digital breadcrumbs wherever they go, both in the real world and online, and most of the people are careless about it. Statistics show how the situation is alarming. Privacy Rights Clearinghouse reports 8,891 data breaches made public since 2005, corresponding to over 11.239 billion records breached <sup>4</sup> ([Privacy Rights Clearinghouse, 2018](#)). In the U.S.A., according to the Bureau of Justice Statistics, in 2014 an estimated 17.6 million individuals experienced any identity theft; this figure is similar to the one reported in 2012, which stood at 12.6 million thefts ([Harrell, 2015](#)). So far, the U.S. FTC brought over 500 enforcement actions protecting the privacy of consumer information addressing well-known companies (Google, Facebook, Uber) as well as lesser-known companies (Upromise, Vizio, SQ Capital) ([FTC, 2017](#)).

The increasingly intensive use of ML and AI techniques applied to Big Data, together with a reduction in the cost of assimilation and processing of data, encourages an overabundant and often indiscriminate collection of information on consumers with the aim of understanding, predicting and influence their behavior ([Jin, 2018](#)). If we add that often the processes of acquisition of

---

<sup>4</sup>As on 13 November 2018

information are transparent to the user, it is understandable how this is no longer able to manage the privacy problems and to make rational and informed decisions about it; even traditional tools such as choice and consent no longer provide adequate protection. There are two sources of user privacy uncertainty:

- The presence of information asymmetry on how the data are processed by the counterpart (Acquisti et al., 2015), and
- The phenomenon identified in the literature as privacy paradox, which indicates the discrepancy (paradoxical, in fact) between what the individual claims to want and her actual behavior. People who claim to care a lot about privacy, usually end up to being inconsistent with their statements, providing personal information with confidence in exchange for discounts or other types of rewards (Spiekermann et al., 2001; Athey et al., 2017).

This behavior does not affect only the naïve but also the most sophisticated individuals; when a consumer has to deal with a decision about privacy, she hardly has all the information she needs to do it. Even when she has all the information, she is unlikely to be able to process them (due to the limited rationality of the individual). Even if she can do it, a series of psychological distortions may intervene. Self-censors themselves when they are aware of the possibility that they are being surveyed, even when knowing they are doing nothing illegal (Stoycheff, 2016).

It is difficult for a consumer to be entirely rational in the face of a decision that involves aspects of privacy. Individuals tend not to protect themselves sufficiently against the perceived privacy risk and to supply excessive amounts of personal data even when they know that they do risk in doing so (Acquisti, 2004). An environment of self-regulation industries does not represent the best solution for the interests of consumers since it does not fill the misalignment of incentives between the parties.

In an environment of self-regulation, giving more information and awareness to users is no longer sufficient to ensure adequate protection. At the same time, on the other hand, whenever consumers are required to make additional effort to protect their privacy (or if this comes at the cost of less smooth user experience), they tend to abandon the technology that would offer them greater protection (Athey et al., 2017).

The topics related to data privacy are also present in the model proposed in this thesis. The variable  $r_{ML}$  was defined as “the future utility that the consumer expects to yield in the future.” This definition implies that, if the consumer’s expectations make her lean towards a more pessimistic view (for example, if many data breaches occur when the consumer decides whether to share data or not), the variable of interest could take on value null or even harmful. In other words, the consumer is induced not to share her information as she expects that third parties will misuse those.

### 3.5.1 Taxonomy of data

A finding of the model proposed states that data are a source of competitive advantage for the firm since they give pricing power and are a proxy for the product’s higher quality. However, model’s finding is about quantity, not about the type of data necessary to create value. To understand which ones are more valuable it is necessary to classify them.

Not all the data collected by the firm are valuable for the firm: provide a taxonomy of data may help to distinguish those who create value from those who do not create value (so to identify those that the company should get rid of). According to the way personal data are acquired, they can be classified as (Schwab et al., 2011):

**Volunteered data** the persistent ones, like name, credit card number, and so on

**Observed data** the dynamic one, like the purchase history of a user



**Inferred data** those derived from the conjoint analysis of volunteered and observed data.

Moreover, having the mere access to data is not like being able to exploit them for commercial purposes or preclude them from others. Data can be classified according to their role for the company, namely:

**Product** as in the case of commercially available databases

**Input** as raw material to improve product functionality and usefulness

**Noncommercial asset** that is byproduct useless for any commercial purpose

Data can also be categorized according to their sensitivity, which can be defined as the degree of importance that a subject gives so some knowledge about herself that, if disclosed, may result in losses of some kind. The definition allows to discriminate:

**Public information** which is information that is a matter of public record

**Private information** that can be used to identify an individual

**Personal information** which is the information belonging to private life, like the details of the domestic life, that cannot be used to identify an individual.

According to the first classification, the firm is interested in all three types. In the second case, the firm looks at data as both a product and an input. In the last case, the firm is more interested in personal information, since they allow it to offer a higher degree of personalization. To the best of the author's knowledge, there is not a definition or personal information measuring unit. The regulator should create one and apply it because if the firm or the consumer are left to define one, they will hardly find common ground (the company has technical knowledge of how the data is used, the consumer could give more value to its information).

### 3.5.2 Problem of identifiability

The figure 3.3 presents an example of different data about an individual ordered according to their sensitivity. When consumers share personal information, there is the direct risk that someone will learn information that the user wished to keep private. One other risk associated with data sharing is identification, undesirable even when just indirect. Combinations of attributes are called a quasi-identifier to differentiate them from directly identifying information like social security number (for example, the combination of 5-digit zip code, birthdate, and gender is a quasi-identifier). Personal preferences like those expressed to many recommender systems may also turn out to be a quasi-identifier, especially if people express unique preferences (Lam et al., 2006).




Figure 3.3. Example of data classification according to their sensitivity.

There is a point in which information sensitivity and the desire of the firm to grasp that information collide. When the consumer does not have the tools to set boundaries, the regulator is required to step in and define a boundary for the firms not to be crossed.

In 2017, 120 Countries adopted laws that outline these boundaries and restrict data collection, either by public or private entities. The remaining Countries have not yet adopted specific laws or fall short in meeting some minimum criteria (Greenleaf, 2017). From a consumer’s perspective, it is somehow unfair that companies operating on a global scale, applying to their users the same contingencies regardless of geography and wealth, are not subject to uniform regulation that ensures to anyone a level of protection at least adequate. In this spirit, there is a call for international organizations to act towards the towards achieving equal rights in data privacy for everyone.

Generally speaking, more data potentially increase recommendation accuracy but also increases the risk of unwanted exposure of personal information. The ideal balance corresponds to a good recommender that does not ask for too much information about customers (Lam et al., 2006). Few would object to improved personalization if it meant, for example, that the barista at a major coffee chain knew a patron’s preferences. However, most would object if the cashier at the local supermarket or whether a recent prescription was effective.

The use of big data in AI-powered products is controversial; some work perfectly with aggregate, anonymous data (e.g., transportation), but many others require individual-specific predictions, making data necessary at an individual level (e.g., healthcare). It is then difficult trying to generalize how to manage consumers data; instead, each case should be assessed independently. Additionally, the application of the laws is complicated in case the company has not directly “stolen” personal data from the user, but has been able to infer them through the AI. Is it still possible to blame the company for violating privacy? How can the company defend itself against these accusations? What if consumers are demanding both algorithmic utility and privacy? How is it possible to protect the consumer against the risk of inferring? These questions require, once again, great regret on the part of the regulator, who must promptly intervene with a legal framework capable of combining and protecting the interests of all the parties involved.

### 3.5.3 Data as a consumer policy law matter

Academics discuss whether Big Data is an antitrust or a consumer protection law matter. Ohlhausen and Okuliar proposed a three-part framework for dealing with Big Data concerns. First, they focus on the nature of the harm, either commercial personal or otherwise. They argue that antitrust should prevail over consumer protection law when there is harm to consumer welfare. Second, they discuss the nature of the consumer-data collector relationship, and they conclude that issues arising from the bargain between these actors are more likely to be a matter of consumer protection law than antitrust. In their last point, they consider the available remedies and related efficiency in resolving particular violations (Ohlhausen and Okuliar, 2015).

Sokol and Comerford (2015) report that Big Data is an antitrust matter only if they are source of unfair competitive advantage, hence by harm consumers; some argued that this could happen if Big Data lead to (1) loss of quality and Innovation, (2) privacy harm, (3) data-driven mergers, (4) perceived strength of scale, network effects and barriers to entry. Their literature review suggests, however, that antitrust law is not suitable to deal with Big Data and their use.

### 3.5.4 Recommendations for the policymaker

As an asset for the firm, data possess unmatched complexity, velocity, and global reach. However, the patchwork of solutions for collecting and using personal data fall short in providing a comprehensive framework to protect the customer. Win-win outcomes can, in contrast, come from creating mutually supportive incentives aimed at stabilizing the personal data ecosystem in a way that creates value for everyone.



Asking antitrust to restrict data collection to those strictly necessary for the operation of the service may seem too conservative and an obstacle to innovation, as the company would be limited in potential solutions that create value from currently unexploited data. It is preferable that such solutions would be put in force by consumer protection laws, to ensure that the consumer is protected from overexposure towards companies anyway. The effort of regulators should be about aligning key stakeholders (people, private firms and the public sector) in support of one another (Schwab et al., 2011):

- innovate around user-centricity and trust,
- define global principles for using and sharing personal data,
- strengthen the dialog between regulators and the private sector,
- focus on interoperability and open standards, and
- continually share knowledge.

The reassuring thing is that no company in the world today (not even security agencies) have access to all a person's data, and even if it did, it would not know how to turn it into a digital twin of an individual. Concerning the model proposed, it is reasonable to suppose that  $U_{ML}$  can never reach its upper bound: the company never sees the bigger picture as a consequence of the fact that every platform used by the user takes care only to collect some information about the user and not all. An e-commerce platform will indeed collect very different data from a virtual personal training platform. In doing so, however, the company cannot leverage the full potential made available by the consumer and therefore she also pays for the costs indirectly associated to lower utility.

Suppose that is possible to let a company take all an individual's data; it would be able to learn a detailed model of the person, “[i]t would surely be a wonderful tool for introspection, like looking at yourself in the mirror, but it would be a digital mirror that showed not just your looks but all things observable about you...” (Domingos, 2015). The digital twin could, however, end up in creating some filter that steers people's life and let them see only what they are expected to like, and anything else. There would be no room for serendipity, for the pleasure of discovery that fascinates and motivates people to live a life of research.

In the end, it is ok if algorithms are not perfect, so that they can introduce something that is a bit of an odd choice and somehow let consumers gain a higher payoff from it. There is indeed something risky and wrong in assuming that everything in the future will be exactly like in the past, and it is also an assumption that hardly fits to how a person's life resembles. Given their conservative focus on past choices, recommendations could overpower consumers into foreseen patterns of consumption and deny them of their capability to evolve, or at the very least reduce the likelihood of radical changes in their tastes.

## Chapter 4

# Conclusions and further developments

ML and smart products are emerging fast, raising concerns about how they can influence the power balance between businesses and consumers. It is necessary, for policymakers, not to lag behind and support a positive adoption of these technologies in the interests of all the stakeholders involved.

This thesis addresses many aspects related to the economy of Machine Learning. The primary goal of this work was to foster debate about the implications of smart systems powered by Machine Learning technologies over the consumer, in particular at microeconomic- and policymaking- level. At the center of the research, a microeconomic model of customer-firm interaction, whereas the product is ML-powered so that its quality improves over time. The firm and the representative consumer are assumed to operate in two interconnected markets, one for the product itself and one for the data that the consumer shares with the firm. The assumption of market interconnection results in a dependence on the product's price and quantity over the data that the firm collects on the secondary market. Following, having chosen two representative companies known to the most for their massive implementation of ML technologies has made it possible to discuss the model to a less-abstract level and verify its applicability; this discussion is merely qualitative. Hence a quantitative, in-depth analysis is required to understand the suitability of the model. The two case studies have later been used to spark a debate on policymakers-related issues, whose future efforts cannot prescribe the investigation of AI technologies and their applications.

The analysis presented should be intended as preliminary, and therefore is not exempt from approximations and inaccuracies to focus on in future research works. The reader should also consider that there may be considerations of critical importance that the author failed to take into account, thereby invalidating some of the conclusions reported. The proposed model is intentionally generic because the spectrum of application cases is varied. It does not mean that future extensions, after adaptation, are more detailed and therefore better in describing a specific case. However, this thesis still offers interesting points of analysis for future research avenues, addressed both to refinement and extension of the proposed model and its empirical validation. Some of them are described in the following section of this chapter.

### 4.1 Future developments of the model

A first step towards the applicability of the model to real situations foresees that some of the parameters hypothesized in the model are estimated empirically. This step is especially crucial for the parameter that incorporates the complementarity between data and product ( $\beta$ ), which is expected to be variable from market to market, if not even from product to product. The same holds for the variable that incorporates the user's expectations about the utility deriving from

sharing personal data with the ML platform ( $r_{ML}$ ), which is not endogenous to the model and therefore requires to be derived (or fixed) a priori.

A further evolution concerns a more in-depth study of the dynamics of net benefit deriving from ML-features. In this thesis, the trend of the variable  $c_{BIAS}$  is assumed symmetric to that of the variable  $U_{ML}$ , and that both had an exponential trend. The author did not provide any empirical evidence to the reader that this is the only possible scenario, hence future versions of the model could ground on the hypothesis that  $c_{BIAS}$  and  $U_{ML}$  have a differently-shaped trend.

Another interesting point to be refined is about investments. In this model, they are limited to the development of a smart algorithm, but they should also account for other IT-related issues, like security countermeasures required for data privacy or the set-up of a proper data center in which store all the consumers' information.

Here, the author considers bias as a variable that affects all consumers in the same way. An extension of the model could consider consumers not all the same, but heterogeneous in the perception of bias (with its consequences on total utility). Plus, it is auspicious to test different versions of the model in which consumers have a different sensitivity to bias; for this variant, it is necessary to modify the incidence of the variable  $c_{BIAS}$  in the utility function of the consumer. Finally, one could study how the model changes as the consumer are entitled to override the biased recommendation of the ML. If the deviation cost is not prohibitive for the consumer, this could be incentivized to choose an alternative and to make the recommendations useless. In this case, the company would be less incentivized to include bias on its platform.

The proposed model foresees the obligation for the consumer to share information with the firm so that he can use its services. If it is possible to share data no longer, free-riding mechanisms are established, whereby the consumer exploits data made available by others without sharing their own, possibly at the cost of obtaining a less personalized experience. The model could also be refined by introducing this aspect, to evaluate the behavior of all consumers compared to that proposed in this thesis.

Finally, rigorous modeling of network effects taking place between the two markets is necessary, to understand how these affect the attraction of new consumers and the quality of the ML feature.

## 4.2 Future research based on the model

In the proposed model, the firm employs the data collected from the user to provide a higher quality product. In reality, the myriad of data that the company collects on users has a different nature depending on the nature of the product or service offered. However, supply chain optimization is another challenging and exciting field for ML to be applied, transforming the efficiency with which many businesses are now operating, turning data into dollars. Cutting-edge solutions can be used to roster staff, improve internal or external logistics, predictive maintenance and so on. An alternative version of the proposed model could contemplate the dependence of the “marginal cost” variable from the data pool available to the company; the model obtained could, once again, be used to draw the behavior of the company concerning the consumer.

A new version of the model could take account for multisided markets in which consumers are not directly charged any price but are, for instance, subject to advertising. It is indeed a very realistic scenario: Alphabet Inc. (formerly known as Google Inc.), for instance, offers many AI-enabled products which are free of charge for the user since the firm profits come from advertising.

Some of the conclusions contained in this work can be used to study how the oligopolistic competition of firms take place in the two interconnected markets, and what is the role of the consumer in this interplay. It is possible to imagine two scenarios, one more unsophisticated in which a firm can exclude the other firms from accessing the consumers' data and a more realistic in which the customer can provide the same data to all the competing firms.

A further expansion of the model could take into account two factors that, in this thesis, have been assumed irrelevant. The decay of data value over time (research shows that the value of data may be transitory or relevant for just a short period (Schepp and Wambach, 2015; Sokol and Comerford, 2016)) and the possibility that the ML system can be attacked by inserting data intentionally biased to perturb the user experience. The scenario can be described introducing a dynamic investment aimed at dampening or eliminating the effect of induced bias.

It has been argued that firms gain market power by excluding others from accessing those data. An ideal extension of this model would make limit the single firm's capability to store and collect information about its users; this role would be taken a public body, that would keep data closed but would also be capable of granting supervised access to the data pool to firms. This solution would move competition on dimensions different from consumer data, for instance, the technologies used to process them. Among other things, this kind of solution would prevent users to lose control of their information and avoid that these could end the wrong hands. In practice, this solution would not be free from privacy concerns, since it would give access to personal information of almost any citizen to a centralized entity <sup>1</sup>. Even if the author agrees that this solution is rather unfeasible and somewhat unethical if this scheme would be welfare-increasing if compared to the setting proposed by the model of chapter 2 (which to some extent reflects how currently the data market works), policy implications can be drawn.

---

<sup>1</sup>In the U.S.A., for example, the Fourth Amendment of the Constitution limit government ability to access and acquire personal belongings, data as in this case.



# Bibliography

- Acemoglu, D. and Restrepo, P. (2016). The race between machine and man: Implications of technology for growth, factor shares and employment. Technical report, National Bureau of Economic Research.
- Acemoglu, D. and Restrepo, P. (2018). Modeling automation. Technical report, National Bureau of Economic Research.
- Acquisti, A. (2004). Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 21–29. ACM.
- Acquisti, A., Brandimarte, L., and Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221):509–514.
- Adomavicius, G., Bockstedt, J. C., Curley, S. P., and Zhang, J. (2013). Do recommender systems manipulate consumer preferences? a study of anchoring effects. *Information Systems Research*, 24(4):956–975.
- Agenzia per l’Italia Digitale (2018). L’intelligenza artificiale al servizio del cittadino [italian]. Retrieved from <https://ia.italia.it/assets/librobianco.pdf>.
- Amazon, Inc. (2018). What data does amazon collect and use? Retrieved from <https://www.amazon.co.uk/gp/help/customer/display.html?nodeId=G6RZ4RMNMLUQRLY2>.
- Amper (2018). Amper music. Retrieved from <https://www.ampermusic.com/>.
- André, Q., Carmon, Z., Wertenbroch, K., Crum, A., Frank, D., Goldstein, W., Huber, J., Van Boven, L., Weber, B., and Yang, H. (2018). Consumer choice and autonomy in the age of artificial intelligence and big data. *Customer Needs and Solutions*, 5(1-2):28–37.
- Andrews, D., Criscuolo, C., Gal, P. N., et al. (2016). The best versus the rest: the global productivity slowdown, divergence across firms and the role of public policy. Technical report, OECD Publishing.
- Athey, S., Catalini, C., and Tucker, C. (2017). The digital privacy paradox: small money, small costs, small talk. Technical report, National Bureau of Economic Research.
- Atomico (2017). The state of european tech 2017. Retrieved from <https://2017.stateofeuropeantech.com/>.
- Barnes, Julian, E. and Chin, J. (2018). The new arms race in ai. Retrieved from <https://www.wsj.com/articles/the-new-arms-race-in-ai-1520009261>.
- Barnhill, A. (2014). What is manipulation. *Manipulation: Theory and practice*, 50:72.
- Baweja, B., Donovan, P., Haeefe, M., Siddiqi, L., and Smiles, S. (2016). Extreme automation and connectivity: The global, regional, and investment implications of the fourth industrial revolution. ubs white paper for the world economic forum annual meeting 2016. *UBS Group AG, Zurich*.

- Belleflamme, P. and Peitz, M. (2015). *Industrial organization: markets and strategies*. Cambridge University Press.
- Blumenstock, J. E. (2016). Fighting poverty with data. *Science*, 353(6301):753–754.
- Botti, S. and Hsee, C. K. (2010). Dazed and confused by choice: How the temporal costs of choice freedom lead to undesirable outcomes. *Organizational Behavior and Human Decision Processes*, 112(2):161–171.
- Bourreau, M. and Gaudin, G. (2018). Streaming platform and strategic recommendation bias.
- Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies ‘engines of growth’? *Journal of econometrics*, 65(1):83–108.
- Brynjolfsson, E., Hu, Y., and Simester, D. (2011). Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science*, 57(8):1373–1386.
- Brynjolfsson, E., Rock, D., and Syverson, C. (2017). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *Economics of Artificial Intelligence*. University of Chicago Press.
- Burda, Y., Edwards, H., Pathak, D., Storkey, A., Darrell, T., and Efros, A. A. (2018). Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*.
- Burguet, R., Caminal, R., and Ellman, M. (2015). In google we trust? *International Journal of Industrial Organization*, 39:44–55.
- Carmon, Z., Wertenbroch, K., and Zeelenberg, M. (2003). Option attachment: When deliberating makes choosing feel like losing. *Journal of Consumer research*, 30(1):15–29.
- Casadesus-Masanell, R. and Halaburda, H. (2014). When does a platform create value by limiting choice? *Journal of Economics & Management Strategy*, 23(2):259–293.
- CB Insights (2016). Artificial intelligence explodes: New deal activity record for ai startups. Retrieved 6th September 2018 from <https://www.cbinsights.com/research/artificial-intelligence-funding-trends/#funding>.
- Chernev, A., Böckenholt, U., and Goodman, J. (2015). Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology*, 25(2):333–358.
- China’s State Council (2017). A next generation artificial intelligence development plan [pdf file]. Retrieved from <https://www.newamerica.org/documents/1959/translation-fulltext-8.1.17.pdf>.
- Cho, J., Lee, K., Shin, E., Choy, G., and Do, S. (2015). How much data is needed to train a medical image deep learning system to achieve necessary high accuracy? *arXiv preprint arXiv:1511.06348*. Retrieved from <https://arxiv.org/pdf/1511.06348.pdf>.
- Christman, J. (2008). Autonomy in moral and political philosophy. *Stanford encyclopedia of philosophy*.
- Chui, M., Manyika, J., Miremadi, M., Henke, N., Chung, R., Nel, P., and Malhotra, S. (2018). Notes from the ai frontier. insights from hundred of use cases. *McKinsey Global Institute*.
- Churchill, O. (2018). China’s ai dreams. *Nature*, 553(7688):S10. Retrieved from <https://www.nature.com/articles/d41586-018-00539-y>.
- Connelly, B. L., Certo, S. T., Ireland, R. D., and Reutzel, C. R. (2011). Signaling theory: A review and assessment. *Journal of management*, 37(1):39–67.

- Cosley, D., Lam, S. K., Albert, I., Konstan, J. A., and Riedl, J. (2003). Is seeing believing?: how recommender system interfaces affect users' opinions. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 585–592. ACM.
- Dalton, R., Mallow, C., and Kruglewicz, S. (2015). Disruption ahead. Retrieved from <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/about-deloitte/us-ibm-watson-client.pdf>.
- Darwin Geo-Pricing (2018). Price for profit with the world's leading dynamic pricing solution for geo-targeted price optimization. Retrieved from <https://www.darwinpricing.com/en/>.
- Datta, H., Knox, G., and Bronnenberg, B. J. (2017). Changing their tune: How consumers' adoption of online streaming affects music consumption and discovery. *Marketing Science*, 37(1):5–21.
- Devlin, H. (2018). London hospitals to replace doctors and nurses with ai for some tasks. Retrieved from <https://www.theguardian.com/society/2018/may/21/london-hospitals-to-replace-doctors-and-nurses-with-ai-for-some-tasks>.
- Directorate for Computer and Information Science and Engineering (2017). Cise funding [pdf file]. Retrieved 6th September 2018 from [http://www.nsf.gov/about/budget/fy2017/pdf/18\\_fy2017.pdf](http://www.nsf.gov/about/budget/fy2017/pdf/18_fy2017.pdf).
- Domingos, P. (2015). *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. Penguin Books Limited.
- Duhigg, C. (2012). Retrieved from <https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>.
- Duprey, R. (2018). Amazon is using big data to dominate the competition. Retrieved from <http://uk.businessinsider.com/amazon-using-big-data-to-dominate-the-competition-2018-2?IR=T>.
- Dutton, T. (2018). An overview of national ai strategies [updated]. Retrieved on the 4th, December, 2018 from <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>.
- England, R. (2018). Chinese school uses facial recognition to make kids pay attention. Retrieved from [https://www.engadget.com/amp/2018/05/17/chinese-school-facial-recognition-kids-attention/?\\_\\_twitter\\_impression=true](https://www.engadget.com/amp/2018/05/17/chinese-school-facial-recognition-kids-attention/?__twitter_impression=true).
- Faden, R. R. and Beauchamp, T. L. (1986). *A history and theory of informed consent*. Oxford University Press.
- Forni, A. (2017). Ai gives customers a valuable resource: Time. Retrieved from <https://www.gartner.com/smarterwithgartner/ai-gives-customers-a-valuable-resource-time/>.
- Frey, C. B. and Osborne, M. A. (2017). The future of employment: how susceptible are jobs to computerisation? *Technological forecasting and social change*, 114:254–280.
- FTC (2017). Privacy & data security update: 2017. Retrieved from [https://www.ftc.gov/system/files/documents/reports/privacy-data-security-update-2017-overview-commissions-enforcement-policy-initiatives-consumers/privacy\\_and\\_data\\_security\\_update\\_2017.pdf](https://www.ftc.gov/system/files/documents/reports/privacy-data-security-update-2017-overview-commissions-enforcement-policy-initiatives-consumers/privacy_and_data_security_update_2017.pdf).
- Galeon, D. and Houser, K. (2017). An ai completed 360,000 hours of finance work in just seconds. Retrieved from <https://futurism.com/an-ai-completed-360000-hours-of-finance-work-in-just-seconds/>.
- General Electrics (2018). Predix platform - digital twin. Retrieved from <https://www.ge.com/digital/predix/digital-twin>.



- Gerbert, P., Hecker, M., Steinhäuser, S., and Ruwolt, P. (2017). Putting artificial intelligence to work. Retrieved from <https://www.bcg.com/publications/2017/technology-digital-strategy-putting-artificial-intelligence-work.aspx>.
- Gillham, J., Rimmington, L., Dance, H., Verweij, G., Rao, A., Roberts, Kate, B., and Paich, M. (2018). The macroeconomic impact of artificial intelligence. Retrieved from <https://www.pwc.co.uk/economic-services/assets/macroeconomic-impact-of-ai-technical-report-feb-18.pdf>.
- Greenleaf, G. (2017). Global data privacy laws 2017: 120 national data privacy laws, including indonesia and turkey.
- Greenstein, S. and Zhu, F. (2012). Is wikipedia biased? *American Economic Review*, 102(3):343–48.
- Gunning, D. (2018). Explainable artificial intelligence (xai). Retrieved from <https://www.darpa.mil/program/explainable-artificial-intelligence>.
- Hall, W. and Presenti, J. (2017). Growing the artificial intelligence industry in the uk. Retrieved from [https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment\\_data/file/652097/Growing\\_the\\_artificial\\_intelligence\\_industry\\_in\\_the\\_UK.pdf](https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf).
- Hallevy, G. (2010). The criminal liability of artificial intelligence entities—from science fiction to legal social control. *Akron Intell. Prop. J.*, 4:171.
- Harrell, E. (2015). 17.6 million u.s. residents experienced identity theft in 2014. Retrieved from <https://www.bjs.gov/content/pub/press/vit14pr.cfm>.
- Hatzius, J., Pandl, Z., Phillips, A., Mericle, D., Pashtan, E., Struyven, D., Reichgott, K., and Thakkar, A. (2016). Productivity paradox v2.0 revisited. *US Economics Analyst*.
- Häubl, G. and Murray, K. B. (2006). Double agents: assessing the role of electronic product recommendation systems. *Sloan Management Review*, 47(3):8–12.
- Hinnosaar, M., Hinnosaar, T., Kummer, M. E., and Slivko, O. (2017). Wikipedia matters.
- Hoyer, W. D. (1984). An examination of consumer decision making for a common repeat purchase product. *Journal of consumer research*, 11(3):822–829.
- International Organization for Standardization (2017). Iso/iec jtc 1/sc 42. Retrieved from <https://www.iso.org/committee/6794475.html>.
- Jain, S. (2018). The potential impact of ai in the middle east. Retrieved from <https://www.pwc.com/m1/en/publications/documents/economic-potential-ai-middle-east.pdf>.
- Jin, G. Z. (2018). Artificial intelligence and consumer privacy. Technical report, National Bureau of Economic Research.
- Kai-Fu, L. (2018). How ai can save our humanity — kai-fu lee. Retrieved from [https://youtu.be/ajGgd9Ld-Wc?list=PLb3\\_87cANkbAE3FQTC7wkHxFtREiGeDbC](https://youtu.be/ajGgd9Ld-Wc?list=PLb3_87cANkbAE3FQTC7wkHxFtREiGeDbC).
- Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., and Fei-Fei, L. (2014). Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732. Retrieved from [https://www.cv-foundation.org/openaccess/content\\_cvpr\\_2014/papers/Karpathy\\_Large-scale\\_Video\\_Classification\\_2014\\_CVPR\\_paper.pdf](https://www.cv-foundation.org/openaccess/content_cvpr_2014/papers/Karpathy_Large-scale_Video_Classification_2014_CVPR_paper.pdf).
- Kingston, J. K. (2016). Artificial intelligence and legal liability. In *International Conference on Innovative Techniques and Applications of Artificial Intelligence*, pages 269–279. Springer.

- Knight, W. (2017). China's ai awakening. *MIT Technology Review*. Retrieved from <https://www.technologyreview.com/s/609038/chinas-ai-awakening/>.
- Komiak, S. Y. and Benbasat, I. (2006). The effects of personalization and familiarity on trust and adoption of recommendation agents. *MIS quarterly*, pages 941–960.
- Korinek, A. and Stiglitz, J. E. (2017). Artificial intelligence and its implications for income distribution and unemployment. Technical report, National Bureau of Economic Research.
- Kwak, I. S., Murillo, A. C., Belhumeur, P. N., Kriegman, D. J., and Belongie, S. J. (2013). From bikers to surfers: Visual recognition of urban tribes. In *BMVC*, volume 1, page 2.
- Lam, S., Frankowski, D., and Riedl, J. (2006). Do you trust your recommendations? an exploration of security and privacy issues in recommender systems. *Emerging trends in information and communication security*, pages 14–29.
- Mankiw, N., Canton, P., and Oliveri, A. (2011). *Macroeconomia*. Zanichelli.
- Manyika, J. (2017). 10 imperatives for europe in the age of ai and automation.
- Manyika, J., Chui, M., Miremadi, M., Bughin, J., George, K., Willmott, P., and Dewhurst, M. (2017). A future that works: automation, employment, and productivity. Retrieved from <https://www.mckinsey.com/~media/mckinsey/featured%20insights/Digital%20Disruption/Harnessing%20automation%20for%20a%20future%20that%20works/MGI-A-future-that-works-Executive-summary.ashx>.
- Marr, B. (2018). How much data do we create every day? the mind-blowing stats everyone should read. Retrieved from <https://www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read/#307c8ed760ba>.
- Metz, C. (2018). As china marches forward on a.i., the white house is silent. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/02/12/technology/china-trump-artificial-intelligence.html>.
- Mozur, p. and Markoff, J. (2017). Is china outsmarting america in a.i.? *The New York Times*. Retrieved from <https://www.nytimes.com/2017/05/27/technology/china-us-ai-artificial-intelligence.html>.
- Musically (2016). Spotify discover weekly playlists have streamed nearly 5bn tracks. Retrieved from <https://musically.com/2016/05/25/spotify-discover-weekly-playlists-5bn-tracks/>.
- Mya (2018). Your a.i. recruiting assistant. Retrieved from <https://hiremya.com/>.
- Neota Logic (2018). Expertise automation for law firms. Retrieved from <https://www.neotalogic.com/industry/law-firms/>.
- Nilsson, N. J. (2009). *The quest for artificial intelligence*. Cambridge University Press.
- OECD (2017). Highlights from the oecd science, technology and industry scoreboard 2017 - the digital transformation: Italy. Retrieved from <https://www.oecd.org/italy/sti-scoreboard-2017-italy.pdf>.
- Ohlhausen, M. and Okuliar, A. (2015). Competition, consumer protection, and the right (approach) to privacy.
- Oxford Dictionary (n.d.). *Artificial Intelligence*. Oxford University Press.
- Partnership on AI (2017). Partnership on AI to benefit people and society. Retrieved from <https://www.partnershiponai.org/>.

- Peirson, V., Abel, L., and Tolunay, E. M. (2018). Dank learning: Generating memes using deep neural networks. *arXiv preprint arXiv:1806.04510*. Retrieved from <https://arxiv.org/pdf/1806.04510.pdf>.
- Perez, S. (2016). Microsoft silences its new a.i. bot tay, after twitter users teach it racism [updated]. Retrieved from <https://techcrunch.com/2016/03/24/microsoft-silences-its-new-a-i-bot-tay-after-twitter-users-teach-it-racism/>.
- Privacy Rights Clearinghouse (2018). Data breaches. Retrieved from <https://www.privacyrights.org/data-breaches>.
- Purdy, M. and Daugherty, P. (2016). Why artificial intelligence is the future of growth. Retrieved from [https://www.accenture.com/lv-en/\\_acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth.pdf](https://www.accenture.com/lv-en/_acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth.pdf).
- Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., Liu, P. J., Liu, X., Marcus, J., Sun, M., et al. (2018). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1(1):18. Retrieved from <https://arxiv.org/abs/1801.07860>.
- Rao, Anand, S. and Verweij, G. (2017). Sizing the prize. Retrieved from <https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html>.
- RapidMathematix (2018). Automated retail pricing. Retrieved from <https://http://rapidmathematix.com/>.
- Rayna, T. and Striukova, L. (2016). 360° business model innovation: Toward an integrated view of business model innovation: An integrated, value-based view of a business model can provide insight into potential areas for business model innovation. *Research-Technology Management*, 59(3):21–28.
- Richter, F. (2018). Netflix continues to grow internationally. Retrieved from <https://www.statista.com/chart/10311/netflix-subscriptions-usa-international/>.
- Russell, S., Dewey, D., and Tegmark, M. (2015). Research priorities for robust and beneficial artificial intelligence. *Ai Magazine*, 36(4):105–114.
- Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of research and development*, 3(3):210–229.
- Saudi Arabia Government (2016). National transformation program 2020. Retrieved from [http://vision2030.gov.sa/sites/default/files/NTP\\_En.pdf](http://vision2030.gov.sa/sites/default/files/NTP_En.pdf).
- Schaefer, M., Sapi, G., and Lorincz, S. (2018a). The effect of big data on recommendation quality.
- Schaefer, M., Sapi, G., and Lorincz, S. (2018b). The effect of big data on recommendation quality: The example of internet search.
- Scheibehenne, B., Greifeneder, R., and Todd, P. M. (2010). Can there ever be too many options? a meta-analytic review of choice overload. *Journal of Consumer Research*, 37(3):409–425.
- Schepp, N.-P. and Wambach, A. (2015). On big data and its relevance for market power assessment. *Journal of European Competition Law & Practice*, 7(2):120–124.
- Schwab, K., Marcus, A., Oyola, J., Hoffman, W., and Luzi, M. (2011). Personal data: The emergence of a new asset class. In *An Initiative of the World Economic Forum*. World Economic Forum.
- Shieber, J. (2018). The not company is looking to start a food revolution from chile. Retrieved from <https://techcrunch.com/2018/07/28/the-not-company-is-looking-to-start-a-food-revolution-from-chile/>.

- Shih, W., Kaufman, S., and Spinola, D. (2007). Netflix. hbs no.9-607-138. *Harvard Business Review*.
- Shoham, Y., Perrault, R., Brynjolfsson, E., Clark, J., and LeGassick, C. (2017). Artificial intelligence index. Retrieved from <https://aiindex.org/2017-report.pdf>.
- Shugan, S. M. (1980). The cost of thinking. *Journal of consumer Research*, 7(2):99–111.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *Nature*, 550(7676):354. Retrieved from [https://deepmind.com/documents/119/agz\\_unformatted\\_nature.pdf](https://deepmind.com/documents/119/agz_unformatted_nature.pdf).
- Sokol, D. and Comerford, R. (2016). Does antitrust have a role to play in regulating big data?
- Solow, R. M. (1957). Technical change and the aggregate production function. *The review of Economics and Statistics*, 39(3):312–320.
- Solsman, J. E. (2018). Youtube’s ai is the puppet master over most of what you watch. Retrieved from <https://www.cnet.com/news/youtube-ces-2018-neal-mohan/>.
- Spice, B. (2017). Carnegie mellon artificial intelligence beats chinese poker players. Retrieved from <https://www.cmu.edu/news/stories/archives/2017/april/ai-beats-chinese.html>.
- Spiekermann, S., Grossklags, J., and Berendt, B. (2001). E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 38–47. ACM.
- Statista (2018). Importance of curated playlists on streaming music services according to consumers in the united states as of march 2018. Retrieved from <https://www.statista.com/statistics/868658/curated-playlists-streaming-music-services/>.
- Stigler, G. J. (1961). The economics of information. *Journal of political economy*, 69(3):213–225.
- Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., et al. (2016). Artificial intelligence and life in 2030. *One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel*.
- Stoycheff, E. (2016). Under surveillance: Examining facebook’s spiral of silence effects in the wake of nsa internet monitoring. *Journalism & Mass Communication Quarterly*, 93(2):296–311.
- Sunstein, C. (2015). Fifty shades of manipulation.
- The Conference Board (2017). Total economy database - growth accounting and total factor productivity, 1990-2016 [spreadsheet]. Retrieved 8th September 2018 from [https://www.conference-board.org/retrievefile.cfm?filename=TED\\_2\\_NOV20171.xlsx&type=subsite](https://www.conference-board.org/retrievefile.cfm?filename=TED_2_NOV20171.xlsx&type=subsite).
- The Economist (2018). Ai helps mass media get up close and personal. Retrieved from <https://huaweiconnect2018.economist.com/ai-helps-mass-media-get-up-close-and-personal/>.
- The Economist Intelligence Unit (2017). Risks and rewards. Retrieved from [https://perspectives.eiu.com/sites/default/files/Risks\\_and\\_rewards\\_2018.2.7.pdf](https://perspectives.eiu.com/sites/default/files/Risks_and_rewards_2018.2.7.pdf).
- The European Commission (2016). Artificial intelligence, robotics, privacy and data protection. In *Room document for the 38th International Conference of Data Protection and Privacy Commissioners*.
- The European Commission (2018). L’intelligenza artificiale per l’europa. Retrieved from <https://eur-lex.europa.eu/legal-content/IT/TXT/PDF/?uri=CELEX:52018DC0237&from=EN>.

- The White House (2016). Preparing for the future of artificial intelligence. *Executive Office of the President, National Science and Technology Council, Committee on Technology*.
- UAE Government (2018). Uae future - 2021-2030. Retrieved from <https://government.ae/en/more/uae-future/2021-2030>.
- University of Cambridge (2015). Artificially-intelligent robot scientist 'eve' could boost search for new drugs. Retrieved from <https://www.cam.ac.uk/research/news/artificially-intelligent-robot-scientist-eve-could-boost-search-for-new-drugs>.
- Vanderbilt, T. (2013). The science behind the netflix algorithms that decide what you'll watch next. Retrieved from <https://www.wired.com/2013/08/qq-netflix-algorithm/>.
- Villani, C. (2018). For a meaningful artificial intelligence. Retrieved from [https://www.aiforhumanity.fr/pdfs/MissionVillani\\_Report\\_ENG-VF.pdf](https://www.aiforhumanity.fr/pdfs/MissionVillani_Report_ENG-VF.pdf).
- Wang, Y. and Kosinski, M. (2017). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. Retrieved from <https://osf.io/zn79k/>.
- Wilson, H. J., Daugherty, P., and Bianzino, N. (2017). The jobs that artificial intelligence will create. *MIT Sloan Management Review*, 58(4):14. Retrieved from <https://sloanreview.mit.edu/article/will-ai-create-as-many-jobs-as-it-eliminates/>.
- World Economic Forum (2016). The future of jobs. Retrieved from <http://www3.weforum.org/docs/GCR2017-2018/05FullReport/TheGlobalCompetitivenessReport2017%E2%80%93932018.pdf>.
- Xiao, B. and Benbasat, I. (2015). Designing warning messages for detecting biased online product recommendations: An empirical investigation. *Information Systems Research*, 26(4):793–811.
- Xiao, B. and Benbasat, I. (2018). An empirical examination of the influence of biased personalized product recommendations on consumers' decision making outcomes. *Decision Support Systems*, 110:46–57.
- Yoon, V. Y., Hostler, R. E., Guo, Z., and Guimaraes, T. (2013). Assessing the moderating effect of consumer product knowledge and online shopping experience on using recommendation agents for customer loyalty. *Decision Support Systems*, 55(4):883–893.
- Ziegler, C.-N. and Lausen, G. (2009). Making product recommendations more diverse. *IEEE Data Eng. Bull.*, 32(4):23–32.

# Appendix - Mathematical proofs

This appendix contains the mathematical proofs of the main findings reported in chapter 2.

## Derivation of the demand functions - equations 2.3

The utility function 2.1 is derived with respect to  $q$  and  $d$ . It is then imposed the condition  $\nabla U(q, d) = (p, v)$  and the two demand functions are derived by solving the linear system of two equations in two variables.

$$U(q_0, q, d) = (V + U_{ML} - c_{BIAS}) \cdot q + r_{ML} \cdot d - 1/2(\alpha q^2 - 2\beta qd + \alpha d^2) + q_0;$$

$$\begin{cases} \frac{\partial U}{\partial q} = V + U_{ML} - c_{BIAS} - \alpha q + \beta d \\ \frac{\partial U}{\partial d} = r_{ML} + \beta q - \alpha d \end{cases} ; \begin{cases} p(q, d) = V + U_{ML} - c_{BIAS} - \alpha q + \beta d \\ v(q, d) = r_{ML} + \beta q - \alpha d \end{cases} ;$$

$$\begin{cases} q = \frac{V+U_{ML}-c_{BIAS}-p+\beta d}{\alpha} \\ d = \frac{r_{ML}-v+\beta q}{\alpha} \end{cases} ; \begin{cases} q = \frac{V+U_{ML}-c_{BIAS}-p}{\alpha} + \frac{\beta}{\alpha^2}(r_{ML}-v+\beta q) \\ d = \frac{r_{ML}-v+\beta q}{\alpha} \end{cases} ;$$

$$\begin{cases} q \left( \frac{\alpha^2 - \beta^2}{\alpha^2} \right) = \frac{\alpha(V+U_{ML}-c_{BIAS}-p) + \beta(r_{ML}-v)}{\alpha^2} \\ d = \frac{r_{ML}-v+\beta q}{\alpha} \end{cases} ;$$

$$\begin{cases} q(p, v) = \frac{[\alpha(V+U_{ML}-c_{BIAS}-p) + \beta(r_{ML}-v)]}{\alpha^2 - \beta^2} \\ d(p, v) = \frac{[\beta(V+U_{ML}-c_{BIAS}-p) + \alpha(r_{ML}-v)]}{\alpha^2 - \beta^2} \end{cases} \blacksquare$$

## Derivation of the number of the sub-games - section 2.6

Let us start considering the equation 2.9; by definition, the efficiency is maximum when it is equal to one; this condition can be imposed to derive the amount of data corresponding to this scenario:

$$1 = 1 - \lim_{D^T \rightarrow ?} e^{-\eta_{DATA} \eta_{ALG} D^T} \iff \lim_{D^T \rightarrow ?} e^{-\eta_{DATA} \eta_{ALG} D^T} = 0 \iff D^T \rightarrow +\infty$$

This value is imposed in equation 2.10, from which it can be seen that an infinite amount of data can only be obtained when  $N \rightarrow \infty$ .

$$+\infty = \sum_{t=1}^T N_t \cdot d_t \iff N \rightarrow \infty \blacksquare$$

## Derivation of the monopoly solution - equations 2.13

The monopolist profit function for the generic  $i$ -th stage game ( $i \neq 1$ ) is defined as  $\Pi = (p-c)q-dv$ , where marginal cost is assumed to be constant and the other quantities are those defined in section 2.5; let us assume  $\phi = 1/(\alpha^2 - \beta^2)$  to lighten the notation:

$$\begin{cases} q(p, v) = \phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] \\ d(p, v) = \phi[\beta(V + U_{ML} - c_{BIAS} - p) + \alpha(r_{ML} - v)] \\ p(q, d) = V + U_{ML} - c_{BIAS} - \alpha q + \beta d \\ v(q, d) = r_{ML} + \beta q - \alpha d \end{cases}$$

The profit function is written expliciting the quantity functions and leaving the price functions indicated; it is later derived with respect  $p$  and then imposed the condition  $\partial\Pi/\partial p = 0$  to find out the monopoly price:

$$\begin{aligned} \Pi &= (p-c)\phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] - v\phi[\beta(V + U_{ML} - c_{BIAS} - p) + \alpha(r_{ML} - v)]; \\ \frac{\partial\Pi}{\partial p} &= \phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] - \alpha\phi(p-c) + \beta\phi v; \\ 0 &= \phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] - \alpha\phi(p-c) + \beta\phi v; \\ p^M &= \frac{V + U_{ML} - c_{BIAS} + c + \beta/\alpha \cdot r_{ML}}{2} \quad \blacksquare \end{aligned}$$

The profit function is rewritten, this time expliciting the price functions and leaving the quantity functions indicated; it is later derived with respect  $d$  and then imposed the condition  $\partial\Pi/\partial d = 0$  to find out the optimal quantity of data for the monopolist:

$$\begin{aligned} \Pi &= (V + U_{ML} - c_{BIAS} - \alpha q + \beta d)q + (r_{ML} + \beta q - \alpha d)d; \\ \frac{\partial\Pi}{\partial d} &= \beta q + r_{ML} + \beta q - \alpha d - \alpha d; \\ 0 &= 2\beta q + r_{ML} - 2\alpha d; \\ d^M &= \frac{1}{\alpha} \left( \beta q + \frac{r_{ML}}{2} \right) \quad \blacksquare \end{aligned}$$

The equations derived above are used to find the value assigned to the exchanged data and the product demand:

$$\begin{aligned} v_S^M &= r_{ML} + \beta q^M - \alpha d^M; \\ v_S^M &= r_{ML} + \beta q^M - \alpha \cdot \frac{1}{\alpha} \left( \beta q^M + \frac{r_{ML}}{2} \right); \\ v_S^M &= r_{ML} + \beta q^M - \beta q^M - \frac{r_{ML}}{2}; \\ v_S^M &= \frac{r_{ML}}{2} \quad \blacksquare \end{aligned}$$

$$\begin{aligned} q_S^M &= \phi[\alpha(V + U_{ML} - c_{BIAS} - p_S^M) + \beta(r_{ML} - v_S^M)]; \\ q_S^M &= \phi \left[ \alpha \left( V + U_{ML} - c_{BIAS} - \frac{V + U_{ML} - c_{BIAS} + c + \beta/\alpha \cdot r_{ML}}{2} \right) + \beta \left( r_{ML} - \frac{r_{ML}}{2} \right) \right]; \\ q_S^M &= \phi \left[ \frac{\alpha}{2} (V + U_{ML} - c_{BIAS} - c) - \frac{\beta}{2} r_{ML} + \frac{\beta}{2} r_{ML} \right]; \\ q_S^M &= \frac{\alpha\phi}{2} (V + U_{ML} - c_{BIAS} - c) \quad \blacksquare \end{aligned}$$

Finally, the product demand is used to determine the optimal quantity of data exchanged:

$$\begin{aligned} d_S^M &= \frac{1}{\alpha} \left[ \beta q_S^M + \frac{r_{ML}}{2} \right]; \\ d_S^M &= \frac{1}{\alpha} \left[ \beta \frac{\alpha \phi}{2} (V + U_{ML} - c_{BIAS} - c) + \frac{r_{ML}}{2} \right]; \\ d_S^M &= \frac{1}{2} \left[ \frac{r_{ML}}{\alpha} + \beta \phi (V + U_{ML} - c - c_{BIAS}) \right]; \quad \blacksquare \end{aligned}$$

## Derivation of the consumer surplus in case of a monopolist firm - equation 2.18

It is given the demand function  $q(p, v) = \phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)]$ . The maximum and the minimum prices that the consumer is willing to pay for the product are respectively known and equal to  $p_{MAX} = p(q = 0) = V + U_{ML} - c_{BIAS} + \beta d^M$  and  $p_{EQ} = p^M = 1/2(V + U_{ML} + c - c_{BIAS} + \beta/\alpha \cdot r_{ML})$ , so that it is possible to compute the consumer surplus:

$$CS_M = \int_{p_{EQ}}^{p_{MAX}} q dp = (p_{MAX} - p_{EQ}) \phi \left[ \alpha(V + U_{ML} - c_{BIAS}) + \frac{\beta}{2} r_{ML} \right] + \frac{\alpha \phi}{2} (p_{MAX} - p_{EQ})^2;$$

The quantity  $(p_{MAX} - p_{EQ})$  is calculated separately, then substituted in the consumer surplus function.

$$\begin{aligned} (p_{MAX} - p_{EQ}) &= \left[ V + U_{ML} - c_{BIAS} + \frac{\beta}{2\alpha} r_{ML} + \frac{\beta^2 \phi}{2} (V + U_{ML} - c_{BIAS} - c) - \frac{V + U_{ML} + c - c_{BIAS} + \beta/\alpha \cdot r_{ML}}{2} \right] \\ &= \frac{1}{2} [(V + U_{ML} - c_{BIAS} - c) + \beta^2 \phi (V + U_{ML} - c_{BIAS} - c)] [\alpha(V + U_{ML} - c_{BIAS})]; \\ &= \frac{1}{2} (V + U_{ML} - c_{BIAS} - c) (1 + \beta^2 \phi); \end{aligned}$$

$$\begin{aligned} CS_M &= \frac{1}{2} (V + U_{ML} - c_{BIAS} - c) (1 + \beta^2 \phi) \phi \left[ \alpha(V + U_{ML} - c_{BIAS}) + \frac{\beta}{2} r_{ML} \right] + \\ &+ \frac{\alpha \phi}{2} \left[ \frac{1}{2} (V + U_{ML} - c_{BIAS} - c) (1 + \beta^2 \phi) \right]^2 \quad \blacksquare \end{aligned}$$

## Derivation of the welfare maximization solution - equations 2.23

The social optimizer profit function for the generic  $i$ -th stage game ( $i \neq 1$ ) is defined as  $W = CS + (p - c)q - dv$ , where marginal cost is assumed to be constant and the other quantities are those defined in section 2.5. The proof of  $\partial CS / \partial p = -q(p)$  and  $\partial CS / \partial p = -v(d)$  is omitted; let us assume  $\phi = 1/(\alpha^2 - \beta^2)$  to lighten the notation:

$$\begin{cases} q(p, v) = \phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] \\ d(p, v) = \phi[\beta(V + U_{ML} - c_{BIAS} - p) + \alpha(r_{ML} - v)] \\ p(q, d) = V + U_{ML} - c_{BIAS} - \alpha q + \beta d \\ v(q, d) = r_{ML} + \beta q - \alpha d \end{cases}$$



The profit function is written expliciting the quantity functions and leaving the price functions indicated; it is later derived with respect  $p$  and then imposed the condition  $\partial W/\partial p = 0$  to find out the welfare-maximizing price:

$$\begin{aligned}
 W &= CS + (p - c)\phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] - v\phi[\beta(V + U_{ML} - c_{BIAS} - p) + \alpha(r_{ML} - v)]; \\
 \frac{\partial W}{\partial p} &= -\phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] + \phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] - \alpha\phi(p - c) + \beta\phi v; \\
 0 &= -\alpha\phi(p - c) + \beta\phi v; \\
 p^W &= c + \frac{\beta}{\alpha}v \quad \blacksquare
 \end{aligned}$$

The welfare function is rewritten, this time expliciting the price functions and leaving the quantity functions indicated; it is later derived with respect  $d$  and then imposed the condition  $\partial W/\partial d = 0$  to find out the optimal quantity of data for the social optimizer:

$$\begin{aligned}
 W &= CS + (V + U_{ML} - c_{BIAS} - \alpha q + \beta d)q + (r_{ML} + \beta q - \alpha d)d; \\
 \frac{\partial W}{\partial d} &= -r_{ML} - \beta q + \alpha d + \beta q + r_{ML} + \beta q - \alpha d - \alpha d; \\
 0 &= \beta q - \alpha d; \\
 d^W &= \frac{\beta}{\alpha}q^W \quad \blacksquare
 \end{aligned}$$

The equations derived above are used to find the value assigned to the exchanged data and the product demand:

$$\begin{aligned}
 v_S^W &= r_{ML} + \beta q^W - \alpha d^W; \\
 v_S^W &= r_{ML} + \beta q^W - \alpha \cdot \frac{\beta}{\alpha}q^W; \\
 v_S^W &= r_{ML} \quad \blacksquare
 \end{aligned}$$

$$\begin{aligned}
 q_S^W &= \phi[\alpha(V + U_{ML} - c_{BIAS} - p_S^W) + \beta(r_{ML} - v_S^W)]; \\
 q_S^W &= \phi\left[\alpha\left(V + U_{ML} - c_{BIAS} - c - \frac{\beta}{\alpha}r_{ML}\right) + \beta(r_{ML} - r_{ML})\right]; \\
 q_S^W &= \phi[\alpha(V + U_{ML} - c_{BIAS} - c) - \beta r_{ML}] \quad \blacksquare
 \end{aligned}$$

Finally, the product demand is used to determine the optimal quantity of data exchanged:

$$\begin{aligned}
 d_S^W &= \frac{\beta}{\alpha}q_S^W; \\
 d_S^W &= \frac{\beta}{\alpha}\phi[\alpha(V + U_{ML} - c_{BIAS} - c) - \beta r_{ML}]; \\
 d_S^W &= \beta\phi\left[V + U_{ML} - c_{BIAS} - c - \frac{\beta}{\alpha}r_{ML}\right]; \quad \blacksquare
 \end{aligned}$$

## Derivation of the welfare maximization solution - equations 2.29

The social optimizer profit function for the generic  $i$ -th stage game ( $i \neq 1$ ) is defined as  $W = CS + k[(p - c)q - dv]$ , where marginal cost is assumed to be constant and the other quantities are those defined in section 2.5. The proof of  $\partial CS/\partial p = -q(p)$  and  $\partial CS/\partial p = -v(d)$  is omitted; let us assume  $\phi = 1/(\alpha^2 - \beta^2)$  to lighten the notation:

$$\begin{cases} q(p, v) = \phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] \\ d(p, v) = \phi[\beta(V + U_{ML} - c_{BIAS} - p) + \alpha(r_{ML} - v)] \\ p(q, d) = V + U_{ML} - c_{BIAS} - \alpha q + \beta d \\ v(q, d) = r_{ML} + \beta q - \alpha d \end{cases}$$

The profit function is written expliciting the quantity functions and leaving the price functions indicated; it is later derived with respect  $p$  and then imposed the condition  $\partial W/\partial p = 0$  to find out the welfare-maximizing price:

$$\begin{aligned} W &= CS + k\{(p - c)\phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)] - v\phi[\beta(V + U_{ML} - c_{BIAS} - p) + \alpha(r_{ML} - v)]\}; \\ \frac{\partial W}{\partial p} &= (k - 1)[\alpha(V + U_{ML} - c_{BIAS}) + \beta r_{ML}] - (k - 1)\alpha p - (k - 1)\beta v - k\alpha p + k\alpha c + k\beta v; \\ 0 &= (k - 1)[\alpha(V + U_{ML} - c_{BIAS}) + \beta r_{ML}] - (2k - 1)\alpha p + k\alpha c + \beta v; \\ \alpha p(2k - 1) &= (k - 1)[\alpha(V + U_{ML} - c_{BIAS}) + \beta r_{ML}] + k\alpha c + \beta v; \\ p_k^W &= \frac{k - 1}{2k - 1} \left( V + U_{ML} - c_{BIAS} + \frac{\beta}{\alpha} r_{ML} \right) + \frac{k}{2k - 1} c + \frac{1}{2k - 1} \frac{\beta}{\alpha} v \quad \blacksquare \end{aligned}$$

The welfare function is rewritten, this time expliciting the price functions and leaving the quantity functions indicated; it is later derived with respect  $d$  and then imposed the condition  $\partial W/\partial d = 0$  to find out the optimal quantity of data for the social optimizer:

$$\begin{aligned} W &= CS + k\{(V + U_{ML} - c_{BIAS} - \alpha q + \beta d)q + (r_{ML} + \beta q - \alpha d)d\}; \\ \frac{\partial W}{\partial d} &= -r_{ML} - \beta q + \alpha d + k\beta q + k r_{ML} + k\beta q - k\alpha d - k\alpha d; \\ 0 &= (k - 1)r_{ML} + (2k - 1)\beta q - (2k - 1)\alpha d; \\ (2k - 1)\alpha d &= (2k - 1)\beta q + (k - 1)r_{ML}; \\ d_k^W &= \frac{k - 1}{2k - 1} \frac{r_{ML}}{\alpha} + \frac{\beta}{\alpha} q^W \quad \blacksquare \end{aligned}$$

The equations derived above are used to find the value assigned to the exchanged data and the product demand:

$$\begin{aligned} v_k^W &= r_{ML} + \beta q^W - \alpha d^W; \\ v_k^W &= r_{ML} + \beta q^W - \alpha \left[ \frac{k - 1}{2k - 1} \frac{r_{ML}}{\alpha} + \frac{\beta}{\alpha} q^W \right]; \\ v_k^W &= r_{ML} - \frac{k - 1}{2k - 1} r_{ML} \quad \blacksquare \end{aligned}$$

$$\begin{aligned}
 q_k^W &= \phi[\alpha(V + U_{ML} - c_{BIAS} - p_k^W) + \beta(r_{ML} - v_k^M)]; \\
 q_k^W &= \phi \left\{ \alpha \left[ V + U_{ML} - c_{BIAS} - \left( \frac{k-1}{2k-1} \left( V + U_{ML} - c_{BIAS} + \frac{\beta}{\alpha} r_{ML} \right) + \frac{k}{2k-1} c + \frac{k}{(2k-1)^2} r_{ML} \right) \right] + \frac{k-1}{2k-1} \beta r_{ML} \right\}; \\
 q_k^W &= \phi \left\{ \alpha \left[ (V + U_{ML} - c_{BIAS}) \left( 1 - \frac{k-1}{2k-1} \right) - \frac{k}{2k-1} c - \frac{k-1}{2k-1} \frac{\beta}{\alpha} r_{ML} + \frac{k}{(2k-1)^2} \frac{\beta}{\alpha} r_{ML} \right] + \frac{k-1}{2k-1} \beta r_{ML} \right\}; \\
 q_k^W &= \frac{\phi}{2k-1} \left\{ k\alpha(V + U_{ML} - c_{BIAS} - c) + \left[ (1-k) \frac{k}{2k-1} \right] \beta r_{ML} + \frac{k-1}{2k-1} \beta r_{ML} \right\}; \\
 q_k^W &= \frac{\phi}{2k-1} \left\{ k\alpha(V + U_{ML} - c_{BIAS} - c) + \frac{(1-k)(2k-1) + k + k-1}{2k-1} \beta r_{ML} \right\}; \\
 q_k^W &= \frac{\phi}{2k-1} \left\{ k\alpha(V + U_{ML} - c_{BIAS} - c) + \frac{(1-k)(2k-1) + k + k-1}{2k-1} \beta r_{ML} \right\}; \\
 q_k^W &= \frac{\phi}{2k-1} \left\{ k\alpha(V + U_{ML} - c_{BIAS} - c) + \frac{-2k^2 + 5k - 2}{2k-1} \beta r_{ML} \right\} \quad \blacksquare
 \end{aligned}$$

Finally, the data value is used to determine the product price and the product demand is used to determine the optimal quantity of data exchanged:

$$\begin{aligned}
 p_k^W &= \frac{k-1}{2k-1} \left( V + U_{ML} - c_{BIAS} + \frac{\beta}{\alpha} r_{ML} \right) + \frac{k}{2k-1} c + \frac{1}{2k-1} \frac{\beta}{\alpha} v_k^W; \\
 p_k^W &= \frac{k-1}{2k-1} \left( V + U_{ML} - c_{BIAS} + \frac{\beta}{\alpha} r_{ML} \right) + \frac{k}{2k-1} c + \frac{1}{2k-1} \frac{\beta}{\alpha} \left( r_{ML} - \frac{k-1}{2k-1} r_{ML} \right); \\
 p_k^W &= \frac{k-1}{2k-1} \left( V + U_{ML} - c_{BIAS} + \frac{\beta}{\alpha} r_{ML} \right) + \frac{k}{2k-1} c + \frac{k}{(2k-1)^2} \frac{\beta}{\alpha} r_{ML} \quad \blacksquare
 \end{aligned}$$

$$\begin{aligned}
 d_k^W &= \frac{k-1}{2k-1} \frac{r_{ML}}{\alpha} + \frac{\beta}{\alpha} \left\{ \frac{\phi}{2k-1} \left[ k\alpha(V + U_{ML} - c_{BIAS} - c) + \frac{-2k^2 + 5k - 2}{2k-1} \beta r_{ML} \right] \right\}; \\
 d_k^W &= \left( \frac{k-1}{2k-1} \frac{1}{\alpha} - \frac{2k^2 - 5k + 2}{2k-1} \frac{\beta^2}{\alpha} \right) r_{ML} + \frac{k\alpha\phi}{2k-1} (V + U_{ML} - c_{BIAS} - c); \\
 d_k^W &= [k-1 - \beta^2(2k^2 - 5k + 2)] \frac{r_{ML}}{\alpha(2k-1)} + \frac{k\alpha\phi}{2k-1} (V + U_{ML} - c_{BIAS} - c) \quad \blacksquare
 \end{aligned}$$

## Derivation of the consumer surplus in case of a social optimizer firm - equation 2.18

It is given the demand function  $q(p, v) = \phi[\alpha(V + U_{ML} - c_{BIAS} - p) + \beta(r_{ML} - v)]$ . The maximum and the minimum prices that the consumer is willing to pay for the product are respectively known and equal to  $p_{MAX} = p(q = 0) = V + U_{ML} - c_{BIAS} + \beta d^W$  and  $p_{EQ} = p^W = c + \beta/\alpha \cdot r_{ML}$ , so that it is possible to compute the consumer surplus:

$$CS_W = \int_{p_{EQ}}^{p_{MAX}} q dp = (p_{MAX} - p_{EQ}) \phi [\alpha(V + U_{ML} - c_{BIAS}) + \beta(r_{ML} - v^M)] + \frac{\alpha\phi}{2} (p_{MAX} - p_{EQ})^2;$$

The quantity  $(p_{MAX} - p_{EQ})$  is calculated separately, then substituted in the consumer surplus function.

$$\begin{aligned}(p_{MAX} - p_{EQ}) &= \left[ V + U_{ML} - c_{BIAS} + \beta^2 \phi (V + U_{ML} - c_{BIAS} - c - \frac{\beta}{\alpha} r_{ML}) - c - \frac{\beta}{\alpha} r_{ML} \right] \cdot \\ &= \left( V + U_{ML} - c_{BIAS} - c - \frac{\beta}{\alpha} r_{ML} \right) (1 + \beta^2 \phi); \end{aligned}$$

$$\begin{aligned}CS_W &= \left( V + U_{ML} - c_{BIAS} - c - \frac{\beta}{\alpha} r_{ML} \right) (1 + \beta^2 \phi) \alpha \phi (V + U_{ML} - c_{BIAS}) + \\ &+ \frac{\alpha \phi}{2} \left[ \left( V + U_{ML} - c_{BIAS} - c - \frac{\beta}{\alpha} r_{ML} \right) (1 + \beta^2 \phi) \right]^2 \quad \blacksquare\end{aligned}$$

## Derivation of the preference conditions of social optimizer versus monopolist solution

Below, the mathematical derivation of the conditions that make the social optimizer firm case better than the monopolist firm case; the condition about the data value requires no calculation and is therefore omitted.

### Price comparison - equation 2.30

$$\begin{aligned}p_S^M &> p_S^W; \\ \frac{1}{2} \left( V + U_{ML} + c - c_{BIAS} + \frac{\beta}{\alpha} r_{ML} \right) &> c + \frac{\beta}{\alpha} r_{ML}; \\ V + U_{ML} + c - c_{BIAS} + \frac{\beta}{\alpha} r_{ML} &> 2c + 2\frac{\beta}{\alpha} r_{ML}; \\ V + U_{ML} - c - c_{BIAS} + \frac{\beta}{\alpha} r_{ML} &> 0; \\ r_{ML} &> -\frac{\alpha}{\beta} (V + U_{ML} - c - c_{BIAS}).\end{aligned}$$

### Quantity comparison - equation 2.31

$$\begin{aligned}q_S^M &< q_S^W; \\ \frac{\alpha \phi}{2} (V + U_{ML} - c - c_{BIAS}) &< \phi [\alpha (V + U_{ML} - c - c_{BIAS}) - \beta r_{ML}]; \\ \alpha (V + U_{ML} - c - c_{BIAS}) &< 2\alpha (V + U_{ML} - c - c_{BIAS}) - 2\beta r_{ML}; \\ 2\beta r_{ML} &< V + U_{ML} - c - c_{BIAS}; \\ r_{ML} &< \frac{V + U_{ML} - c - c_{BIAS}}{2\beta}.\end{aligned}$$

**Shared data - equation 2.32**

$$\begin{aligned}
d_S^M &> d_S^W; \\
\frac{1}{2} \left[ \frac{r_{ML}}{\alpha} + \beta\phi(V + U_{ML} - c - c_{BIAS}) \right] &> \beta\phi(V + U_{ML} - c - c_{BIAS} - \beta r_{ML}); \\
\frac{r_{ML}}{\alpha} + \beta\phi(V + U_{ML} - c - c_{BIAS}) &> 2\beta\phi(V + U_{ML} - c - c_{BIAS}) - 2\beta^2\phi r_{ML}; \\
\frac{r_{ML}}{\alpha\beta\phi} &> V + U_{ML} - c - c_{BIAS} - 2\frac{\beta}{\alpha}r_{ML}; \\
r_{ML} \left( \frac{1}{\alpha\beta\phi} + 2\frac{\beta}{\alpha} \right) &> V + U_{ML} - c - c_{BIAS}; \\
r_{ML} \left( \frac{1 + 2\beta^2\phi}{\alpha\beta\phi} \right) &> V + U_{ML} - c - c_{BIAS}; \\
r_{ML} &> (V + U_{ML} - c - c_{BIAS}) \frac{\alpha\beta\phi}{1 + 2\beta^2\phi}.
\end{aligned}$$