



European
University
Institute

DEPARTMENT
OF LAW

WORKING PAPERS

LAW 2018/11
Department of Law

Consumer law and artificial intelligence
Challenges to the EU consumer law and policy stemming
from the business' use of artificial intelligence

Final report of the ARTSY project

Agnieszka Jabłonowska
Maciej Kuziemski
Anna Maria Nowak
Hans-W. Micklitz
Przemysław Pałka
Giovanni Sartor

European University Institute
Department of Law

**CONSUMER LAW AND ARTIFICIAL INTELLIGENCE
CHALLENGES TO THE EU CONSUMER LAW AND POLICY
STEMMING FROM THE BUSINESS' USE OF ARTIFICIAL
INTELLIGENCE**

Final report of the ARTSY project

Agnieszka Jabłonowska
Maciej Kuziemski
Anna Maria Nowak
Hans-W. Micklitz
Przemysław Pałka
Giovanni Sartor

EUI Working Paper **LAW** 2018/11

This text may be downloaded for personal research purposes only. Any additional reproduction for other purposes, whether in hard copy or electronically, requires the consent of the authors. If cited or quoted, reference should be made to the full names of the authors, the title, the working paper or other series, the year, and the publisher.

ISSN 1725-6739

© Agnieszka Jabłonowska, Maciej Kuziemski, Anna Maria Nowak,
Hans-W. Micklitz, Przemysław Pałka, Giovanni Sartor 2018
Printed in Italy
European University Institute
Badia Fiesolana
I-50014 San Domenico di Fiesole (FI)
Italy
www.eui.eu

Abstract

Potential regulation of use of artificial intelligence by business should minimize the risks for consumers and the society without impeding the possible benefits. To do so, we argue, the legal reaction should be grounded in an empirical analysis and proceed case-by-case, bottom-up, as a series of responses to concrete research questions. The ambition of this report has been to commence and facilitate that process. We extensively document and evaluate the market practice of the corporate use of AI, map the scholarly debates about (consumer) law and artificial intelligence, and present a list of twenty five research questions which, in our opinion, require attention of regulators and academia.

The report is divided into four sections. The first explains our understanding of the concepts of “artificial intelligence” (a set of socio-technological practices enabled by machine learning and big data) and “consumer law” (various legal instruments concretizing the principles of the weaker party protection, non-discrimination, regulated autonomy and consumer privacy). The second section documents the ways in which the business uses artificial intelligence in seven sectors of the economy: finance and insurance, information services, energy and “smart solutions”, retail, autonomous vehicles, healthcare and legal services. For each analyzed sector we study the gains for the businesses stemming from the deployment of AI, the potential gains, but also challenges for consumers, as well as third party effects. In the third section, we repeat the analysis through the lens of four general “uses” of AI by businesses in various sectors: knowledge generation, automated decision making, advertising and other commercial practices and personal digital assistants. Finally, in the fourth section, we present the questions which we believe should be addressed in the next stage of the research. We cluster them into: normative questions about regulatory goals, technological and governance questions about regulatory means, and theoretical questions about concepts and preconceptions.

Keywords

Consumer law and policy, challenges, artificial intelligence, business use

Author contact details

Agnieszka Jablonowska

Faculty of Law and Administration
University of Lodz

AJablonowska@wpia.uni.lodz.pl

Maciej Kuziemski

School of Transnational Governance
European University Institute
Florence, Italy

Maciej.Kuziemski@eui.eu

Anna Maria Nowak

Law Department
European University Institute
Florence, Italy

Anna.Maria.Nowak@eui.eu

Hans-W. Micklitz

Law Department
European University Institute
Florence, Italy

Hans.Micklitz@eui.eu

Przemysław Palka

Law Department
European University Institute
Florence, Italy
Yale Law School
New Haven, CT, United States

Przemyslaw.Palka@eui.eu

Giovanni Sartor

Law Department
European University Institute
Florence, Italy

Giovanni.Sartor@eui.eu

Table of contents

INTRODUCTION: WHAT ARE WE DOING HERE AND WHY? 1

1. SETTING THE SCENE 4

 1.1. Artificial intelligence in context..... 4

 1.2. The bird’s eye view: lawyers reflecting upon artificial intelligence 7

 1.3. EU consumer law: What’s that all about? 8

 1.3.1. Protection of the weaker party..... 9

 1.3.2. Regulated autonomy 12

 1.3.3. Non-discrimination..... 14

 1.3.4. Consumer privacy..... 18

2. SECTOR ANALYSIS: WHAT DOES BUSINESS USE AI FOR, WHAT ARE THE GAINS AND THE RISKS?..... 19

 2.1. Finance and insurance 20

 2.2. Information services 22

 2.3. Energy and “smart solutions” 25

 2.4. Retail 27

 2.5. Autonomous vehicles 29

 2.6. Healthcare..... 30

 2.7. Legal..... 33

3. ANALYSIS BY USE: DATA GATHERING AND KNOWLEDGE GENERATION, AUTOMATIC DECISIONS, ADVERTISING, PERSONAL ASSISTANTS 35

 3.1. Data gathering and knowledge generation 35

 3.2. Automated decision making 38

 3.3. Advertising and other commercial practices 48

 3.4. Personal Assistants (“Algorithmic Consumers”, “Cyber-Butlers”) 51

4. TIME FOR PHASE TWO: FURTHER QUESTIONS IN PLACE OF CONCLUSIONS..... 57

 4.1. Normative questions (goals)..... 57

 4.2. Technological (means) 64

 4.3. Governance (means)..... 65

 4.4. Theoretical/conceptual 68

LIST OF LEGAL ACTS 71

BIBLIOGRAPHY 73

Introduction: What are we doing here and why?¹

The overall objectives of the ARTSY project have been to 1) map and document the use of artificial intelligence and other algorithmic techniques by business, 2) analyze and evaluate this practice through the prism of rules and principles of EU consumer law, including personal data protection law, 3) reflect on the challenges to consumer law and policy stemming from this practice; 4) prepare the next stage of the research, in which normative scientists, engineers and governance experts will collaborate. This will be a longer project, which will inquire into possible regulatory goals and means of achieving these goals, and put forward concrete policy recommendations. All this should be executed with a scholarly mindset, in dialogue with other researchers studying related subjects, hence 5) our ambition to also extensively map the state of the art in (consumer) law and artificial intelligence scholarship.

The structure of this project rests on four methodological premises. First, that the normative (legal) research in the field of emerging technologies must be **based on empirical studies**. This way, we analyze the actual practices (what is really going on), instead of merely an idea of what artificial intelligence is and what it is used for. That is why we begin with extensive mapping and documenting. Second, such research must be **grounded in concrete normative theory**, in order to provide a clear threshold for evaluation. That is why we have chosen EU consumer law. This is not to say that this is the only normative threshold, or the best normative threshold. On the contrary, many other perspectives – including human rights – are needed, but to move forward in a dialogue, we need these perspectives to be clear and specified. Third, **there will be no one-size-fits-all, top-down solution to the “AI challenge”, because there is no one “AI challenge”**. Different technologies, used by different types of actors, for different purposes, pose different risks and might need very different solutions. That is why we aim at identifying as many concrete challenges as possible, and asking specific questions. **Fourth**, to regulate AI, or even to ponder whether to regulate it, a **wise interdisciplinary approach** is needed. Lawyers (and other normative scientists) are well-trained to provide an evaluation of phenomena, and to propose **goals**. However, they are not the only ones out there who should be proposing means for achieving those goals. That is where expertise of engineers (concerning **technical means**) and governance experts (concerning **political means**) is needed. Therefore we limit our normative analysis to evaluation and asking questions about policy goals, but leave the questions about means for the next stage of the research.

This project has been a response to several overlapping phenomena. On the one hand, **market practice**. Businesses and different corporate actors currently employ artificial intelligence and other algorithmic techniques in essentially all sectors of the economy, for tasks ranging from pattern recognition and behavior prediction, to optimization and influence on consumers. Our research has shown that companies use AI in all the studied sectors: finance and insurance, information services, energy and “smart solutions”, retail, autonomous vehicles, healthcare and legal. Across those sectors, AI is used to generate new knowledge, automate the decision-making processes, create and deliver targeted advertising and other commercial practices, and power digital agents (chatbots, personal

¹ The authors would like to thank the participants of the workshop *Before Machines Consume the Consumers. Consumer Law and Policy in the Algorithmic Age*, organized on 28-29 June 2018 at the European University Institute in Florence, for their feedback, suggestions and comments, which made this report much better thought-through. In particular, they would like to thank Professor Urs Gasser for his invaluable support at all stages of the project, and express their gratitude towards the Berkman Klein Center for Internet and Society at Harvard University, for funding which made the research leading to this report possible.

digital assistants). These uses, which influence the economic decisions of consumers, as well as their autonomy, privacy and other protected values are sufficient reason to at least ask the question: is there a role for consumer law to play in these dynamics?

All this created a sense of **political urgency**. With numerous relevant events taking place on a weekly basis, reports being published faster than one human being can read them, and within the narrative of “we have to do something now!”², what is more important than ever is a **cold legal analysis that is critical by design**. Scholars agree that the development of AI has a range of impacts in settings as diverse as job provision (Frey et al., 2013), human rights protection (Eubanks, 2018), healthcare delivery (Jiang et al., 2017), and energy management (Lösch and Schneider, 2016). Yet the evident lack of clarity as to the expected magnitude of disruptions, does not preclude a sense of political urgency shared by decision-makers, scientists, and entrepreneurs. Considered against the backdrop of environmental challenges, unsustainability of growth models, or democratic decay, one may – and many have – come to conclusion that we are living through a critical juncture – a brief period of time with a heightened probability that choices will affect outcomes, launching new pathways and transforming institutions (Capoccia, 2015). In such times, it is particularly relevant to scrutinize actors, processes and dynamics, while options are still wide open (Jasanoff and Kim, 2015). To critically observe and ponder, where indeed we need to act fast, and where it is rather preferable to stand and stare.

Finally, all this gave rise to a specific **scholarly practice**. Numerous published articles treat artificial intelligence as one technology, posing a uniform set of problems across different sectors and consequently calling for a uniform set of solutions. Questions about “liability for artificial intelligence systems”, “personality for artificial agents” or “ways of dealing with the black-boxing problem” are not uncommon, overlooking the fact that often the term “artificial intelligence” is used as an umbrella category, capturing extremely diverse types of socio-technological phenomena. On the other extreme, there are publications studying very specific issues, overlooking the context in which they are embedded, or works already done in other “fields”, which however can often be useful directly or indirectly. We term this scholarly practice **Phase One of AI & Law scholarship**. This stage, necessary but insufficient, can be characterized by high excitement, a sense of urgency, initial familiarization with the technical and social knowledge, testing initial intuitions and hypotheses. We claim it is now time for **Phase Two: moving to concrete questions, addressed in a wise interdisciplinary dialogue, aware of the broader context and other debates, though proceeding not top-down, but bottom-up, case by case, problem by problem**.

The project aims at making several contributions and serving several objectives. First, it wants to provide a **map – description and analysis** – of the market practice and of the scholarship. We want this document to serve as a reference point for anyone – scholars, policymakers, journalists, activists – who might feel overwhelmed with the sheer amount of material they need to deal with. Second, we propose a **methodological contribution**, suggesting that while researching challenges posed by AI, one should proceed in steps: empirical analysis, normative evaluation, proposing goals, technical means and governance means, all this within the overarching frame of a wider theoretical reflection. Third, having executed the mapping and explained the method, we make all the material **freely available to any research groups** who might want to pick up on what we did here, and actually proceed to stage two.

² For an overview of national and transnational AI strategies, see: <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>

A couple of reservations. This was a short, preparatory, **sprint** research project, by design aimed at preparing avenues for further work. We do not aspire to **solve** anything here. On the contrary, our claim is that the debates, both scholarly and policy-oriented, would benefit a lot from **slowing down**, not jumping straight to regulatory conclusions, and devoting slightly more time to trying to **understand** what is going on first. Moreover, the ambition for his project was to be **comprehensive in subjects**, i.e. to provide an overview of spaces where in-depth research into consumer law & AI could/should be done. With time and person-power limitations, we needed to make choices, and we chose breadth over depth. This follows from our conviction that one should study **concrete questions**, but one needs to have a general idea of the whole problem to pose them in the first place. Finally, this document is a bit experimental in form. In the times when everyone writes and so there is no time to read, we wanted to make it more **readable** and **enjoyable** to study. Hence, sections 2 and 3 consist of introductions and tables, which we hope facilitate familiarization with the material.

All this said, we would welcome **any** feedback with great joy, especially when it is critical. The task at hand is a hard one. Let us take it seriously.

1. Setting the scene

1.1. Artificial intelligence in context

Beginning a paper by asserting that no commonly accepted definition of artificial intelligence/robotics exists seems to be a good practice within the field. Having done so indirectly, we would like to argue that this is not really a problem.

Definitions, roughly speaking, perform two functions: 1) clarify the meaning of words (what are we referring to when employing the terms “artificial intelligence”, “robotics” and “algorithmic techniques”?) or 2) increase our understanding of the entities that the words refer to (assuming we all refer to the same things when using the term “algorithm”, what are the characteristics of these things?).³ Regarding the first: as we argue, “artificial intelligence” is *not* exactly the same *thing* in autonomous weapons as in high frequency trading and it neither *does* the same *thing* in personalized advertising as in self-driving cars: it poses different problems depending on who uses it, in what way, and for what purpose. As a consequence, there will be no Artificial Intelligence Act, regulating the whole phenomenon top-down, administered by the Artificial Intelligence Agency (just like there is no Internet Act, or Internet Agency). **Hence, there is no need for one, all-encompassing definition**, at least for the purpose of legal research and regulation. And even sector-wise, as long as we know what we are talking about, definitions are not necessary. Ever wondered how is it possible that the legal system manages to regulate the behavior of human beings, even though no commonly accepted definition of a “human being” exists?

As long as we know what we are talking about. But do we? That is the second function of definitions. Again, part of the problem with the law and artificial intelligence scholarship is that our understanding of AI is often influenced by linguistic intuitions and sci-fi movies and books, rather than by the socio-economic reality itself. However, increasing knowledge and understanding does not have to come through *the* definition, or even a series of definitions. It can come in various forms, like explanations, exemplifications or a narrative. A short and concise statement might meet our aesthetic needs – scholars like order and, yes, definitions – but given how complex the phenomenon is, it might lead us to generalizations and a loss of nuance.

This said, being content with not having one all-encompassing definition does not free us from our scholarly obligation to understand what we are talking about when invoking artificial intelligence or algorithms.

When speaking about the “regulation of AI” or “challenges of AI”, or essentially regulation or challenges of any technology, we have to bear in mind that by this “technology” one actually should understand a **socio-technological practice**, i.e. what different actors *do* with it, what the technology *enables* them to do, and what the consequences of these technologically enabled *actions* are. An algorithm for assessing an individual’s willingness to pay, when simply placed on a repository online, does not pose any challenge yet. A company using this algorithm to show lower flight prices to less wealthy consumers and higher prices to the wealthier ones might pose one challenge, while a company using the same algorithm in order to show unaffordable prices to people professing a certain religion definitely poses another one. In the words of Jack Balkin (2015a):

When we consider how a new technology affects law, our focus should not be on what is essential about the technology but on what features of social life the technology makes newly salient. What problems does a new technology place in the foreground that were previously underemphasized or deemed less important?

³ These, again roughly speaking, are respectively the “nominal” and “real” definitions. See also: <https://przemyslaw.technology/2015/11/15/the-trouble-with-what-is-x-questions/>

And further:

What lawyers call “technology” is usually a shorthand for something far more complex. When we talk about “technology,” we are really talking about (1) how people interact with new inventions and (2) how people interact with other people using those new inventions or presupposing those new inventions. What we call the effects of technology are not so much features of things as they are features of social relations that employ those things (...) Innovation in technology is not just innovation of tools and techniques; it may also involve innovation of economic, social, and legal relations. (Balkin, 2015a)

In this sense, what we need to do here is to highlight what techno-sociological practices we are planning to study, and say something about the technology that enables them. If the reader still thinks that a definition of artificial intelligence is in place, we refer to Russell and Norvig (2016) who survey numerous definitions, pointing out that scholars tend to define the discipline differently, depending on whether they are concerned with the process (“thinking”) or outputs (“acting”), and whether they take as a threshold “being like a human”, or being “rational”. To give an example, Kurzweil et al. (1990) define artificial intelligence as “*The art of creating machines that perform functions that require intelligence when performed by people*”, while Haugeland (1985) defines it as: “*The exciting new effort to make computers think . . . machines with minds, in the full and literal sense.*”

This does not get us anywhere. We might spend hours debating the concepts of “rationality”, “intelligence”, “acting like humans”, but that in no way enriches the understanding of the practice nor the challenges it poses to (consumer) law. Fortunately, the approach of the businesses is much more concrete.

What is actually happening now, under the label of “artificial intelligence”, is businesses using machine learning to make sense out of vast amounts of data (big data), to generate new knowledge, and act upon that knowledge in order to optimize certain processes, and undertake new tasks, previously impossible. Now, one by one.

The story usually goes like this (see, for example, Domingos, 2015; Alpaydin, 2016). Once upon a time, to have a computer “do” something, a human programmer needed to write the software explaining the machine how to go about solving some problem. This had certain limitations, given that for many tasks that seem simple for humans, we do not actually know how to formalize them. A classic example: recognizing a photo of a cat. Every person can tell if there is a cat on a picture. But how to “program” an instruction for recognizing a cat, for a machine that only “sees” pixels out there?

Machine learning puts the idea upside down. Instead of telling a machine how to do something, we provide it with a huge amount of input and output data – for example, one billion photos with cats and one billion photos without them – and let it “figure out” how to do this by itself. What the computer gets from humans is the *learning algorithm* (how to learn something) and *data*, but it “programs itself” when it comes to finding a recipe for solving a given problem. This, by the way, is one of the meanings of the “black box” or “inexplicability”: in machine learning, we do not know *how* a computer did something. What we do know, however – what many industries representatives often do not mention openly – is 1) what the task was for which the machine was being trained; 2) what learning algorithm has been employed; 3) what (types of) data it has been trained on.

Data. “Data is the new oil”. Indeed, it has been data, big data, that enabled the machine learning revolution. Numerous tasks (machine translation, image/voice recognition, medical diagnostics) were automatable because we had 1) the learning algorithms; 2) sufficiently high computing power; 3) sufficiently big data sets. How much is sufficient? That highly depends on the nature of the task. But usually a lot. The process of digitization led to the creation of the (big) data sets, as a by-product of simply living. More and more parts of our lives have been moved into the “cyberspace”, using smartphones all the time, shopping online and sharing our lives on social media. And if someone wants to learn something, there is really a lot to be learned there. Consumer data is valuable. Ever wondered

why Facebook and Google are worth billions of dollars – even though they charge you nothing for using their services?

So what we have until now is: companies using machine learning techniques – training machines, using learning algorithms and huge amounts of data – to automate various tasks that were not automatable using traditional programming. If we know what data is, let us maybe take a look at the other widely used term, which is seldom defined: “**algorithm**”. In its report on algorithmic accountability, the World Wide Web Foundation stated:

Although typically defined as a set of “encoded procedures” or “a logical series of steps for organizing and acting on a body of data to quickly achieve a desired outcome”, the term “algorithm” is often intended to describe a larger intersection of code, data and automated decisions. Originating from computer science and used in various social science disciplines, the term has been used to convey various meanings on the intertwining of human and machine decision inputs, and the extent to which the term includes code, data and ecosystems often varies. (World Wide Web Foundation, 2017)

Indeed, also in legal scholarship, the term “algorithm” is often employed in quite a loose way, in order to capture various situations of “computers doing stuff”. In the technical sense, it is usually defined as a “method of solving a problem”, a “formalized set of steps to realize a particular task”, in short: an algorithm is a **recipe** for doing something. This recipe can be **computed** by a human or a machine.

For example, think of an algorithm for establishing a scholar’s h-index.⁴ The instruction (algorithm) would be: order the scholar’s papers in a list, by the number of citations; look for the last paper where the number of citations is equal or greater to position in the ranking; that position number is the h-index. This instruction can be executed by a human or a computer. Google Scholar does this for every author who has a profile there. Here we need a clarification: if an algorithm is just an instruction, what is the name for the “entity” executing it in the sentence “Google Scholar does so for every author”. Google Scholar’s employee? Of course not! It would be tempting to say “Google Scholar’s algorithm”, but the algorithm is just the instruction, not the entity executing it.

There is a word we are lacking. Try to solve the following puzzle: IF humanity THEN human, IF robotics THEN robot, IF artificial intelligence THEN... what, exactly? We seem to be missing a word here. Some call it “artificial agent” (Chopra and White, 2011), some call it “autonomous systems”, some just also say “algorithms”. None of this is perfect. In this report we try to be descriptive, speaking rather of “business using AI-powered tools”. Yet when we need a term, we rely on the good old “artificial agent”.

These agents can be categorized in various ways. By embodiment: robots vs software agents; degree of autonomy: autonomous vs automatic; the task they realize; the legal good they put in danger; the economic good their functioning is supposed to optimize. Here, we want to claim, one should be careful not to fall too deep into the dichotomies – “autonomy” or “embodiment” are much more often different shades of green than pure blue and yellow.

What is more, terms like “algorithms”, “artificial intelligence” or “artificial agents”, as well as the “challenges they pose”, are often meant to relate to much wider sets of socio-technological phenomena. This is what we call the **degree of inseparability**. To a large extent, it is not possible, and probably not even desirable, to separate the phenomena of big data, online platforms, regulation by the code, Internet of Things⁵, internet to begin with, and artificial intelligence. When necessary, we do analyze them jointly.

⁴ A metric equal the number (h) of an author’s publications that have been cited at least h times. For example, an author with two papers cited at least two times will have h-index equal 2, five papers cited at least five times h index equal 5 etc. See: <https://en.wikipedia.org/wiki/H-index>

⁵ The term “Internet of Things” (IoT) has already been used in the literature for more than a decade. In 2005, the International Telecommunication Union described the upcoming development of “new forms of communication between people and things, and between things themselves”. The ITU further asserted: “A new dimension has been added to the world of

Summing up: **When speaking of “artificial intelligence” we refer to a socio-technological practice of companies (or other actors) using machine learning tools to generate computer-readable knowledge out of big amounts of data, and further use that knowledge to optimize certain processes and undertake new types of actions, for example to predict consumer (individual/group) behavior, influence it, take decisions etc. The entities performing these tasks are called artificial agents. Depending on the type of agent or its task, the exact legal challenges and regulatory responses might be very different in content and in form.**

1.2. The bird’s eye view: lawyers reflecting upon artificial intelligence

Artificial intelligence, as a technology and as a phenomenon, has attracted attention from a variety of disciplines, including legal scholarship. Even though it could seem that the subject is new, parts of the legal community have been interested in AI for quite some time already. Indeed, the regular attendees of the bi-annual *International Conference on AI and Law* (first held in Boston in 1987), as well as the regular readers of the *Artificial Intelligence and Law* journal (established in 1992) would probably strongly oppose the claim that AI & Law is a new discipline, or that artificial intelligence is a new phenomenon that has only recently been discovered by the legal community. However, upon a closer examination of what subjects excite the mainstream legal scholarship nowadays, both would probably admit that, indeed, there is a certain divergence in the types of questions that have been asked.

To put it short (and slightly simplistically): for most of its existence, AI & Law as a discipline has been concerned with issues of formalization of legal norms and legal reasoning, essentially struggling with the question: **given the properties of law and information systems, to what extent is it possible to have computers do what lawyers do?** This problem has been approached by scholars working at the intersection of legal theory and informatics. It has remained at the periphery of the mainstream legal scholarship, which arguably either found it too boring, or could not understand all the mathematical formulas used there. An overview of this endeavor is neatly presented by Bench-Capon et al. (2012) in: *A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law*. Numerous publications on the subject have been written, and the discipline still thrives (Ashley, 2017). This is not the strand of literature we wish to review here.⁶

The approach to the intersection of law and artificial intelligence that has been slowly crawling into the mainstream legal scholarship for the past five years, kindling excitement in lawyers and the public, is the policy-oriented approach. Here, to again put it simplistically, the general question is: **given that artificial intelligence is being successfully employed by different actors across numerous sectors of socio-economic life, with all its promises and challenges, what should we do about law?** It is this strand within which the ARTSY project situates itself, and it is this strand of which we would like to provide an overview here.

Artificial intelligence as a subject of (potential/insufficient) regulation attracted attention of scholars from numerous **branches of law**. This includes **criminal law** and the questions of criminal liability or *mens rea* (Hallevy, 2013); **competition law** with the questions of data as market power and algorithmic collusion (Ezrachi and Stucke, 2016); **intellectual property law** with the questions about authorship of computer-generated works and inventions and the possibility to use AI to enforce IP rights (Abbott, 2016), **private law** i.e. tort and contract (Pagallo, 2013) and international law, where extensive debates about the use of automated weapon systems took place (Krishnan, 2016).

information and communication technologies (ICTs): from anytime, anyplace connectivity for anyone, we will now have *connectivity for anything*.” Item identification systems, sensors technologies, embedded intelligence as well as further miniaturization and nanotechnology have been described as key enablers of the IoT development (International Telecommunication Union, 2005; see also: Weber and Weber, 2010).

⁶ This said, an attempt, within this strand precisely, to automate the legal evaluation of consumer contracts and privacy policies is the objective of the other AI & consumer project hosted at the Law Department of the European University Institute: <https://claudette.eui.eu/>.

A distinction cutting across these branches can be drawn between the **law-centric** approaches, where the primary object of inquiry was still the law, in the context of artificial intelligence; and **machine-centric** approaches, where the primary object of study would be the socio-technological phenomenon seen from the point of view of laws goals rather than established institutions. Within the first strand, questions would be asked about the suitability of law for artificial intelligence, for example: is a contract made by a machine binding according to the current doctrine (Allen and Widdison, 1996; Sartor, 2009)? Or who is responsible for a tort committed by a robot as the law stands now (Pagallo, 2013)? Within the second strand, the authors would depart from an issue, a challenge posed by AI-systems, and ask questions about possible ways of solving it, either by interpreting the existing legal norms (Balkin, 2015b), or by proposing new regulatory solutions (Calo, 2015; Scherer, 2015).

Another way to categorize the scholarship is by the type of **legal problem** that the author identifies as central. Numerous contributions have touched upon the questions of **liability** for actions of artificial agents/ AI-powered systems (Asaro, 2016), **transparency** of algorithmic decision making (Perel and Elkin-Koren, 2017), **accountability**, **data bias** and potential for **discrimination**, **personalization** of goods/services/communications, the legal **status** of artificial agents (including calls for potential granting them legal **personality**), privacy and more recently questions about **AI governance**. Within the last strand, numerous solutions have been proposed, including the creation of administrative bodies certifying autonomous systems (Scherer, 2015), “tinkering” the algorithms to deduce the rationale of their functioning (Perel and Elkin-Koren, 2017), granting the companies that employ them a specific fiduciary status (Balkin, 2015b, see also: Lightbourne, 2017), employing various techniques for ex ante and ex post oversight and control of AI-powered systems (Kroll et al., 2017; Kim, 2017a).

We provide a much more detailed analysis of the particular parts of the field in section 3, discussing various perspectives on business’ use of AI. However, what should be made clear now is: moving forward with the legal research about artificial intelligence and law *will* require a case by case approach. Meaning: there will *always* be the question of liability, of transparency, of accountability etc.; but the *answer* to these questions might differ across the sectors/uses/types of systems and/or legal goods endangered.

1.3. EU consumer law: What’s that all about?

We have so far clarified what we mean by “artificial intelligence” and “algorithms”, as well as what the subjects were that the scholarship has found exciting until now. Before moving to the actual market analysis, we should still explain what we mean by the “concrete normative threshold”, and in our case, European consumer law. One way to approach the problem would be to through the existing legal instruments that make up the system of consumer law in the EU. We do so indirectly below. However, since these instruments were certain particularizations of the general ideas – principles – created in a given socio-economic-technical context, it might just as well be that their content and form do not necessarily best suit the challenges of the AI. Hence, what we explain in detail are these *principles* – a much more stable foundation of the EU consumer law (though of course also not carved in stone). These guiding principles, around which the analysis is built, are:

- Protection of the weaker party
- Regulated autonomy
- Non-discrimination
- Consumer privacy

However, before getting into the technical legal analysis, we want to give the reader a “feeling” of what the issues at stake are. Consider the following story:

Feel the Consumer Struggle: A Short Story

Frank was devastated. Rushing out of the sterile building hosting the CoolMed Inc., he was wondering whether he could lose weight by simply walking faster. Cholesterol three times the norm, BMI 39, severely obese. “No progress” – two words uttered by the disappointed doctor still echoed in his ears – “if you want to live past 60, you MUST lose weight”.

It’s not that he didn’t try. He REALLY tried. After hearing the doctor’s recommendation three months ago he bought all these veggies, signed up for the gym, got his bike serviced... Yet in the end most of the veggies rotted in the fridge and the two sessions he attended in the gym turned out to be ridiculously expensive, given the cost of the three-month \$150 membership. With all the deadlines at work, additional jobs he accepted in hope of finally getting the pay-raise, children hitting adolescence in a painful way, and his marriage essentially falling apart, Frank really did not have the time to do all the healthy cooking and grocery shopping. He was also really too tired to go jogging in the morning, after the by-now-unfortunately-regular 5 hours of sleep every night.

With tears in his eyes, Frank left the CoolMed behind him and started walking towards his car. That is when he got approached by the marketer. “Looks like you could use some help!” – uttered a man of his age, though way more fit. “With Lose10poundFor10pounds™ your body will burn the fat four times faster than normal. You don’t even need to do sport, simply walking will do all the burning for you!”. Frank was not convinced, though down there, in the subconscious part of his mind, a faint voice kept shouting: “yes, give me this magic!”. Seeing the doubt on Frank’s face, the marketer changed the tone from euphoric to compassionate and said: “Look, I know this sounds improbable. But I only sell those because it helped me, and I want to help others. This is a photo of me half a year ago” – the marketer showed Frank a picture – “and with Lose10for10, that’s the short name, I managed to go down 30 kilograms in six months. It’s really effective, British lab stuff”. Ok, this was convincing. He had a photo. He had proof.

“How much?” – Frank asked. “\$99 for a 30 days pack” – answered the marketer. “I’ll take three then” – Frank replied. And thought “I hope this will save me. What is 300 bucks when speaking about health”. “What a moron” – the marketer thought. He still could not believe that these cheap tricks they learned during the training, and the photoshopped picture of himself, would convince anyone in their right mind.

Lose10poundFor10pounds™ obviously did not work, being made of chalk that used to stand next to the green tea. It took Mark a few weeks to realize he got tricked, but when he finally did, he got furious. He googled the product, and realized it was a scam. He tried to contact the company, but they pointed out that the contract he signed clearly stated that he accepts the possibility of the product not working, and that there would be no refunds.

Frank lost money, time and nerves.

The purpose of the story above is to give the reader a feeling of what EU consumer law is about. The principles discussed below, which have informed the regulation, remain the same. However, as the reader can certainly imagine by now, the challenges to consumer law have changed in form and in substance. In the next two sections, which form the core of this document’s content, we rely on these principles as our normative threshold, our normative intuition, informing our analysis of what should be perceived as a challenge and why.

1.3.1. Protection of the weaker party

In the Western World, the birthday of consumer law is associated with the declaration of president Kennedy in 1962. It took more or less ten years before the message arrived in Europe. Here it got a particular twist – the European democracies put the emphasis on consumer *protection* law. The founding years were marked by attempts to legitimize and define the consumer as the weaker party in economic transactions. ‘Weak’ was associated with the consumption sphere. The consumer is the one who is not doing business. Scholars made many attempts to draw parallels to the role of the worker in the production process. However, these efforts did not really lead to practical consequences. The doctrinal debate focused on the imbalance of power: whether there is one at all and if so, how the imbalance can

be defined. The idea was generally linked to the perceived asymmetry in the level of knowledge, experience and bargaining power⁷.

From the 1980s onward, the EU – back then still the European Economic Community – gradually took over consumer law. As it is argued below, this shift was not without effect for the understanding of the principle of the weaker party protection in European law. Originally, however, the involvement of the Community in the field of consumer law was driven by similar considerations as those prevalent in the Member States. Even if consumers were no longer perceived as entirely helpless due to the case law of the Court of Justice⁸, the early Community-level debates clearly emphasised the need of their protection. The 1975 Council resolution observed that the new market conditions, such as an increased abundance and complexity of goods and services, have shifted the balance between suppliers and customers in favour of the former. According to the Council, such an accentuated imbalance resulted in the need to keep consumers “better informed of their rights and protected against abuses”.⁹ Shortly after, in the second consumer programme of 1981, further factors were mentioned that transform the relationships between market participants, including the “new data processing and telecommunications technology”¹⁰.

The perception of the consumer as a weaker party in need of *protection* was reflected in the subsequent legislative developments at the Community level, most notably in the principle of minimum harmonisation upon which the common rules had originally been based and the areas in which they were developed (product safety¹¹, unfair terms¹²). Already at this early stage, however, the protection granted to consumers was balanced against market considerations. By way of illustration, instead of banning the practice of selling at the doorstep, the European lawmaker provided consumers with the right to withdraw¹³. The interaction between market and protection is also visible in the concept of legitimate expectations found in the instruments on product liability¹⁴ and, less conspicuously, on unfair terms¹⁵.

Over time, the importance of market-related objectives increased, partially at the expense of the original protective goals. The shift was also reflected at the rhetorical level. The adopted policy documents no longer focused on consumer protection as such, but rather on making it possible for consumers “to realise the benefits of the internal market”¹⁶ and “benefit from the opportunities presented by the information society”¹⁷. Despite Treaty reforms that occurred in the meantime¹⁸, most of the Commission’s legislative initiatives affecting the position of consumers continued to derive their legal basis from the provisions related to the functioning of the internal market and have increasingly followed the full harmonisation approach (Tonner, 2014). The functions of the advancing harmonisation policy were associated not so much with the removal of barriers to trade, but rather with the promotion of

⁷ See cases: C-89/91 *Shearson Lehman Hutton*; C-240/98 to C-244/98 *Océano*; C-168/05 *Claro*.

⁸ See in particular case C-210/96 *Gut Springenheide*, ECLI:EU:C:1998:369 and the average consumer notion.

⁹ Council resolution of 14 April 1975 on a preliminary programme of the European Economic Community for a consumer protection and information policy [1975] OJ C92/1.

¹⁰ Council resolution of 19 May 1981 on a second programme of the European Economic Community for a consumer protection and information policy [1981] OJ C133/1.

¹¹ Directive 92/59/EEC on general product safety (later repealed and replaced by Directive 2001/95/EC) and Directive 85/374/EEC on product liability.

¹² Directive 93/13/EEC on unfair terms in consumer contracts, Article 5.

¹³ Directive 85/577/EEC on contracts negotiated away from business premises; later repealed and replaced by Directive 2011/83/EU on consumer rights.

¹⁴ Directive 85/374/EEC on product liability, Article 6.

¹⁵ See the elaboration on the requirements of good faith in the preamble of Directive 93/13/EEC on unfair terms in consumer contracts.

¹⁶ COM(2002) 208 final.

¹⁷ COM(95) 519 final.

¹⁸ In particular the introduction of the current Article 169 of the Treaty on the functioning of the European Union.

“consumer confidence” in the viability of the border-free market (Weatherill, 2001). Such an approach envisages a consumer who wishes to engage actively in the increasingly integrated market and clearly builds upon the notion of an “average consumer” (see also: Mak, 2013). Under the Internal Market doctrine, consumer protection law turned into consumer law *without* protection (Micklitz, 2012). This view is supported by the type of legal tools applied to protect consumer interests, many of which did not really aim at the protection of the weaker or at social redistribution, but rather at the creation and safeguarding of “access justice” (Micklitz, 2011).

To be sure, even though the principle of the weaker party protection has to some – or even considerable – extent been instrumentalized to further market integration, it continues to occupy a prominent place in the European legal framework. The principle continues to be invoked by the Court of Justice in the different areas of the *acquis*, such as Directive 2005/29/EC on unfair commercial practices (UCPD)¹⁹, Directive 93/13/EEC on unfair terms in consumer contracts (UCTD)²⁰ or Directive 1999/44/EC on consumer sales²¹, to provide consumer-friendly interpretations. This is particularly true when the existential consumer interests are at stake as has been the case in the aftermath of the financial crisis (see also: Grundmann, 2016). In an attempt to better reflect this nuanced picture, particularly with consumer law entering ever broader fields of the economy in the new millennium, the EU started to differentiate between the responsible, the confident and the vulnerable consumer. The latter two, and particularly the vulnerable consumer – even if only gradually entering the European scene (Reich, 2016), come closest to the original concept of the weaker party, the one where the consumer finds him or herself confronted with a big and/or experienced business. The time seems right to ask a further question: does digitalization, and particularly the rise of AI, bring anything new to this picture?

The extent to which consumers should be treated as the weaker parties in the digital market has been the subject of debate. Overall, there are two conflicting positions. The first perceives digital consumers as most likely to actually correspond with the responsible/confident model, which the EU has created for them. Such an interpretation has also found support of some national courts, which consider digital consumers as more technology-savvy, active, well-informed, cautious and attentive²². The second position, by contrast, sees all digital consumers as potentially vulnerable, due to, among others, the technological advantage of their counterparts. This view has been prominently represented by Mik (2016), who criticized the European legal framework, especially the UCPD, for failing to advance the idea of vulnerability as a “dynamic state”. According to the author “each consumer can be vulnerable in its own way and ... vendors have the technological capacity to exploit temporary vulnerabilities – not just those caused by age, mental infirmity or credulity” (Mik, 2016). A similar argument is made by Helberger et al. (2017), who observe that, in view of the possibility for traders to “target individuals”, more legal attention should be paid to the individual characteristics and vulnerabilities (Helberger et al., 2017; see also: Sax et al., in the context of mHealth apps). This thread of discussion is also linked to the potentially new dimensions of information asymmetries in the algorithmic society. First of all, in the age of big data and learning algorithms that are able to identify the needs and risk profiles of particular consumers, it is valid to ask questions about the degree and scope of the informational advantage that controllers of the algorithms possess over consumers. One may argue that such an advantage is no longer limited to the product or service and the overall market experience, but also extends to the characteristics of consumers themselves. By contrast, consumers may be unaware not only of the existence of such factors, but also of the very fact of being subject to the technological influence. What is more: even if they possess such knowledge, they might still be unable to draw meaningful conclusions (e.g. evaluate the quality of the advice provided or the accuracy of an algorithmic prediction made). All this might

¹⁹ Cases C-59/12 *Zentrale zur Bekämpfung unlauteren Wettbewerbs*; C-388/13 *UPC Magyarország*.

²⁰ Case C-137/08 *Pénzügyi Lízing*; C-415/11 *Mohamed Aziz*.

²¹ Cases C-497/13 *Faber*; C-149/15 *Wathelet*.

²² Judgment of the Polish Supreme Court of 17 September 2014 in case I CSK 555/13.

place additional strains on the tools traditionally used to level the asymmetry between consumers and traders, such as the classical disclosure duties.

Overall, the extent to which the weaker party protection is revitalized through the use of big data and learning algorithms by the businesses appears as one of the most pertinent questions to be addressed in the course of further research on consumer law and AI.

1.3.2. Regulated autonomy

The importance of personal autonomy in the discussions about democratic rights and individual freedom is undisputed (Young, 1986). As observed by Reich (2014), “every liberal legal order has the autonomy of private parties as its basic philosophy”. The idea of autonomy is also clearly reflected at the market level. The freedom of suppliers to decide whether and on what conditions to enter markets is considered fundamental to the functioning of open market economies, as is the freedom of the customers to choose the products and services and the conditions on which they are willing to enter into legal relationships. Both of these perspectives (the freedom of active market participants to decide and the freedom of passive market participants to choose) are reflected in the principle of the freedom of contract, the importance of which has been stressed exhaustively in both scholarship and jurisprudence²³.

A detailed discussion of the conception and the role of autonomy within the legal system is a monumental task that goes way beyond the remit of this project. For a concept of such a cardinal importance, autonomy remains a strikingly undefined notion. The discussions carried out by the private lawyers tend to focus upon its reflections in the principle of contractual freedom, which, admittedly, are already difficult enough (Weatherill, 2013). The differences between particular conceptions of the freedom of contract, which developed in different historical moments across the different jurisdictions (see, in particular: Micklitz, 2015; Kennedy, 2006), are indeed remarkable. These relate to the overall importance attached to 1) the individual freedom (as the popular generalization goes: the premise of the common law is that ‘where there is a remedy, there is a right’ while the civil law tradition takes an opposite perspective; Worthington, 2015); 2) the source of an individual obligation in contractual ordering (as reflected in the transformations of the objective theory of contract and the will theory; Kennedy, 2000) and 3) the principal rationale for statutory interventions (e.g. utilitarian or paternalistic; Micklitz, 2015).

These different perspectives are still reflected in the modern discussions on the role of autonomy in the digital age. A good illustration of this point is the distinction made by Mik between a “contractual” and a “regulatory” perspective on the challenges to consumer autonomy in the technology-mediated environments. The former perceives the problem as one of a potentially “defective consent”, with an ensuing focus on the doctrines like misrepresentation, mistake or undue influence. The latter is more concerned with the “systemic dimension” of the analyzed phenomenon, which, in turn, justifies the deployment of particular “regulatory instruments”. Examples mentioned in the article include mandatory disclosure rules and outright prohibitions found in the EU consumer *acquis* (Mik, 2016). Arguably, however, the line between both perspectives is more difficult to draw.

Indeed, the evolution of EU private law has largely been shaped by its focus on market integration. This internal market orientation accounts not only for the regulatory nature of the European *acquis*, but also for its distinct approach to autonomy of the legal subjects (Comparato and Micklitz, 2013; Comparato et al, 2016). Unlike national legal frameworks, private law of the EU is not particularly concerned with the theoretical underpinnings such as private autonomy or the freedom of contract, but rather instrumentalizes them to achieve its objectives (Micklitz and Patterson, 2013). Arguably, however, it does so without undermining the concept of autonomy as such, but rather places this concept in a new perspective. Take one of the basic regulatory tools used by the European legislator – the information

²³ See e.g. Grundmann et al. (2001); Draft Common Frame of Reference rule II. – I:102; case C-283/11 *Sky Österreich*.

duty. Its widespread deployment is based on a premise that, for the integration project to be sustainable, the information asymmetry between consumers and traders needs to be reduced. In other words, for the same market-related purposes, the autonomy of the consumers needs to be re-established through qualified information (see, for example, Article 6 of Directive 2011/83/EU on consumer rights (CRD)). The importance of transparency for furthering autonomy is also visible in the EU's General Data Protection Regulation 2016/679 (GDPR). In this context, particular attention should be drawn to the idea of consent as one of the grounds of the lawful processing of personal data. Pursuant to Article 4(11) of the GDPR consent should be "freely given, specific, informed and unambiguous". The regulation goes on to specify that, in assessing whether consent is freely given, utmost account shall be taken of whether, *inter alia*, the performance of a contract, including the provision of a service, is conditional on consent to the processing of personal data that is not necessary for the performance of that contract (Article 7(4)). Note that, while the GDPR is without doubt a regulatory instrument, its provisions on consent vividly evoke of the "contractual perspective". All in all – irrespective of the criticism of the protection through information (Howells, 2005; Ben-Shahar and Schneider, 2014), the disputed practical relevance of the provision on "tying" data subject's consent (Article 29 Working Party, 2017b) and the thorny problem of consenting to the secondary use of data (Pearce, 2015) – in both cases the tools used by the legislator are aimed to assist in consumer's (data subject's) decision-making without fundamentally affecting the nature of the good or service offered. Admittedly, this enhanced freedom to choose on the demand side is not without (restraining) effect on the suppliers' freedom to decide. In some areas, like in unfair terms law, this bent towards consumer's interests is even clearer (see also Micklitz and Reich, 2014). Most of the time, however, the imposed restraints do not amount to a direct inroad into the content of the bargain or determine which goods or services could be lawfully placed on the market (see also: Grundmann, 2002). The EU legislator tends to perceive them not only as adequate to the level of risk addressed, but also a small price to be paid by the traders in exchange for the benefits derived from market integration.

The latter argument is particularly visible in the field of telecom, energy, transport and partly financial services, i.e. the markets previously dominated by state incumbents which were gradually opened to competition via EU law. The opening of the markets towards new entrants enlarged the autonomy of both the supplier and the consumer. However, the supplier had to accept that the new freedoms are coming at a price – regulated autonomy. Private law of the EU therefore, simultaneously, "establish[es] market freedoms and therefore increase[s] private autonomy" and "set[s] boundaries to this newly created autonomy" or, in other words, "frames" it (Micklitz and Patterson, 2013; Reich, 2014). Such a perception – together with the recognition that different legal norms may affect autonomy of the particular market actors to a different degree – sheds a more nuanced light on the classical debate about the extent to which regulation limits private autonomy. As such, it also allows us to move beyond this predominant context for assessing the role of autonomy in private law and fits perfectly well the European approach on AI.

With this in mind it is worth giving some thought to the instrument which the literature reviewed recognizes as the most promising candidate, besides the GDPR, for addressing the challenges that consumers face in the age of AI – Directive 2005/29/EC on unfair business-to-consumer commercial practices (Mik, 2016; Helberger, 2016; Sax et al., 2018). The act brings yet another example of the rules that seek to increase the consumer's freedom to choose, without exceedingly intervening with the trader's freedom to decide. Indeed, for a practice to be qualified as unfair within the meaning of the UCPD's general clause two conditions need to be met: firstly, the trader must fail to comply with the "requirements of professional diligence" and, secondly, the practice must (at least) be likely to "materially distort" the economic behavior of the average consumer with regard to the product²⁴. This autonomy-based rationale is additionally highlighted in the provisions on "aggressive" practices, where

²⁴ Directive 2005/29/EC on unfair business-to-consumer commercial practices, Article 5(2).

a reference to a “significant impairment of the average consumer’s *freedom of choice or conduct* with regard to the product” is made²⁵.

The noticeable interest of legal scholars in the UCPD seems to reflect the increasingly widespread concerns about the impact of data-intensive business models, and particularly the more and more sophisticated tools of technological influence, on consumer’s autonomy. With its emphasis on consumers’ decision-making powers the UCPD indeed appears to be most apt to provide legal responses. Interestingly, however, as recently highlighted by Sax et al., autonomy in the European fair trading law has not been conceptualized sufficiently. Therefore, even though autonomy constitutes a legal term, it requires a reference point outside the legal system (Sax et al. 2018; Micklitz, 2006). In their recent contribution Sax et al. attempt to identify these reference points in the realm of ethics. All in all, they distinguish between three requirements for autonomy: independence, authenticity and options. The former refers broadly to the capacity of being in control of one’s life, or, more narrowly (referring to Valdman, 2010), to the possibility of exercising “managerial control over the motivating elements of one’s psychology”. According to this reading, an autonomous agent should be able to consider all the available information and options independently and “decide how her values, desires, and goals inform her intentions for acting”. Note that this dimension of autonomy comes close to the discussions about dignity (see also section 3.4 below; Brownsword, 2017). The second requirements of autonomy in the reading of Sax et al., authenticity, refers to the question whether the relevant values, desires and goals can be described as truly personal to the decision-maker. This part of analysis focuses on what it means to identify with one’s own positions and to what extent a person can be manipulated into forming such values, desires and goals. Last but not least, for autonomy of a person to be preserved, an adequate range of options to choose from should be available to him or her. The importance of availability and awareness of available options also growingly transpires from the contributions of other authors investigating the digital transformation from consumers’ perspective (Mik, 2016; Gal and Elkin-Koren, 2017). Further research devoted to consumer’s autonomy in the age of AI is certainly encouraged.

1.3.3. Non-discrimination

The non-discrimination principle reached the consumer law rather late. The origins lay elsewhere. Both in the EU and in the US the principle of non-discrimination – or, somewhat more broadly, the principle of equal treatment – has traditionally been applied to vertical relations (between the citizens and the institutions governed by public law). In the United States its key manifestation is the Equal Protection Clause set out in the Fourteenth Amendment to the United States Constitution as interpreted in a seminal line of the Supreme Court’s case law²⁶. In Europe, the concepts of equality and non-discrimination are generally derived from the Member States’ constitutional traditions, the law of the European Union, and human rights law, in particular the European Convention on Human Rights (McCrudden and Prechal, 2009). Due to this complex and evolving background, including considerable transatlantic influences, a coherent theoretical framework underpinning the European principle of non-discrimination is difficult to decode.

The importance of the principle of equal treatment within the legal order of the Community (now the Union) was confirmed early on by the Court of Justice²⁷. In the light of this case law, the principle is considered as one of the general principles of EU law and requires that comparable situations are not treated differently and that different situations are not treated in the same way, unless such treatment is

²⁵ Directive 2005/29/EC, Article 8.

²⁶ See, in particular judgments of 17 May 1954 *Brown v. Board of Education*, 347 U.S. 483 (1954) and *Bolling v. Sharpe*, 347 U.S. 497 (1954).

²⁷ See, e.g. cases C-117/76 *Ruckdeschel and Others*; C-106/83 *Sermide*; C-309/89 *Codorniu v Council*; C-133/93 *Crispoltoni and Others*; equality and discrimination are also explicitly mentioned in Article 2 of the Treaty on the European Union as core values upon which the EU is founded.

objectively justified (see generally: Tridimas, 2007). This basic understanding of equality underpins its more specific manifestations. In this respect, McCrudden and Prechal distinguish, among others, between equality “as an adjunct to the protection of particularly prized public goods” and equality as non-discrimination on the basis of (or equal treatment irrespective of) specific characteristics (McCrudden and Prechal, 2009).

The former understanding of equal treatment is most directly linked to the nature of the EU (originally the EEC) as a regional political and legal system with a primary aim of market integration. In this respect, equality is mainly associated with the prohibition of discrimination on grounds of nationality and serves as one of the instruments for furthering market-related goals. This dimension of equal treatment is clearly reflected in both primary and secondary law of the EU, most notably in Article 18 of the Treaty on the Functioning of the European Union, which specifically prohibits discrimination on grounds of nationality. Non-discrimination of the nationals of other Member States is also ingrained in the concept of EU citizenship and plays an important role in the free movement law²⁸. References thereto can further be found in Directive 2006/123/EC on services in the internal market and in the more recently adopted Directive 2014/92/EU on payment accounts and Regulation 2018/302 on geo-blocking. With a growing number of specific references to non-discrimination of this kind and the gradual expansion of the EU’s social dimension, the boundaries between the first and the second understanding of equality from the perspective of nationality become blurred. This second understanding corresponds to the conception of equality that has underpinned the development of anti-discrimination laws, which focused on specific protected categories such as racial or ethnic origin, on both sides of the Atlantic. Notably, not only did this dimension of EU anti-discrimination law develop significantly later than in the US (perhaps with the exception of the non-discrimination based on gender), but many of its building blocks are also considered to have drawn from the American ideas (De Búrca, 2012).

Indeed, the principle of non-discrimination on grounds of a specific protected category draws from the developments in the former colonizing countries, which have gradually introduced laws against race discrimination. In the EU the developments were due to an active Court of Justice, which provided the ground for taking political action via secondary EU law, via the amendment of the Treaty and, since 2000, via the Charter of Fundamental Rights. Indeed, particularly from the 2000s onward, a considerable body of secondary EU legislation has been adopted, which gave the principle of non-discrimination based on protected characteristics a more tangible expression and expanded its horizontal meaning (Reich, 2014).

Analysis of the legal framework as well as of the literature concerned with the horizontal dimension of the principle of equal treatment in the EU leads to the following observations. First of all, there appears to be a limited range of *characteristics* on the basis of which discrimination is prohibited by law. Such protected characteristics include, in particular, nationality, gender, racial and ethnic origin, religion and belief, age and sexual orientation, and do not extend to economic grounds (here the rationales of the weaker party protection and the regulated autonomy might be more relevant). At present, the most comprehensive anti-discrimination framework in areas of relevance to the present project concerns racial and ethnic origin²⁹, gender³⁰ and, with all reservations expressed above, nationality (for a comprehensive overview, see: European Union Agency for Fundamental Rights and Council of Europe, 2018).

²⁸ Articles 20, 45 and 49 TFEU; Directive 2004/38/EC on the right of citizens of the Union and their family members to move and reside freely within the territory of the Member States.

²⁹ Directive 2000/43/EC implementing the principle of equal treatment between persons irrespective of racial or ethnic origin and case C-83/14 *CHEZ*.

³⁰ Directive 2004/113/EC implementing the principle of equal treatment between men and women in the access to and supply of goods and services and case C-236/09 *Test Achats*.

Secondly, the balance between the principle of non-discrimination and the competing principles of private law, such as freedom of contract, appears to depend on the *area of law* under analysis (Reich, 2014). Considerations of equality are particularly prominent in labor law, in which the most fully-fledged framework has been established covering all the protected characteristics mentioned above³¹. This tendency reflects the earlier developments in the US where a prohibition of discrimination based on “race, color, religion, sex, or national origin” in the employment context exists since the adoption of Title VII Civil Rights Act of 1964³². In other areas, most notably consumer law, to which the present project is devoted, the principle of non-discrimination applies less coherently. The differences arise as regards both the types of protected characteristics and the types of consumer transactions with respect to which discrimination is prohibited. While some legal acts focus only on one protected characteristic and have a broadly defined subject matter (Directive 2000/43/EC on racial and ethnic origin which prohibits discrimination in social protection, health care, education and access to goods and services, including housing, which are available to the public³³), others follow an exactly opposite logic (Directive 2014/92/EU which prohibits discrimination based on all grounds set out in Article 21 of the Charter with respect to access to payment accounts³⁴). Such an incoherent approach appears to be attributed at least partially to the lack of political consensus on further-reaching harmonization. This is evidenced by the disagreement about a proposed directive on implementing the principle of equal treatment between persons irrespective of religion or belief, disability, age or sexual orientation, which sought to implement the principle of discrimination on these grounds outside the labor market³⁵. Overall, provisions aimed at combating discrimination in consumer markets tend to focus either on 1) goods and services of particular social importance (specified/network services such as payment, insurance and banking, services of general economic interest³⁶); 2) goods and services which are generally available to the public.

Of relevance is finally the *direct or indirect* nature of the discriminatory treatment. The former describes a situation where a person is treated less favorably “on grounds” of a protected characteristic (such as sex or racial or ethnic origin), than another is, has been or would be treated in a comparable situation. The latter refers to a situation where an “apparently neutral provision, criterion or practice” would put persons possessing a given protected characteristic “at a particular disadvantage” compared to persons not possessing that characteristic.

Importantly, while the EU law typically provides for only limited defenses to direct discrimination, indirect discrimination can often be justified by reference to a more general defense. Indeed, pursuant to several anti-discrimination directives, indirect discrimination might be considered lawful, provided that the provision, criterion or practice “is objectively justified by a legitimate aim, and the means of

³¹ Directive 2000/43/EC implementing the principle of equal treatment between persons irrespective of racial or ethnic origin; Directive 2000/78/EC establishing a general framework for equal treatment in employment and occupation; Directive 2004/38/EC on the right of citizens of the Union and their family members to move and reside freely within the territory of the Member States; Directive 2006/54/EC on the implementation of the principle of equal opportunities and equal treatment of men and women in matters of employment and occupation (recast); Directive 2014/54/EU on measures facilitating the exercise of rights conferred on workers in the context of freedom of movement for workers.

³² For an analysis from a perspective of the modern developments such as automated decision-making see: Barocas and Selbst (2017) and Kim (2017).

³³ Discrimination based on sex is prohibited in a similar range of areas. Interestingly, Directive 2004/113/EC also contains a specific provision on insurance, which led to a controversial judgment of the Court of Justice in C-236/09 *Test Achats*.

³⁴ Directive 2014/92/EU on the comparability of fees related to payment accounts, payment account switching and access to payment accounts with basic features, Article 15.

³⁵ COM(2008) 426, cf. some national laws: German General Act on Equal Treatment (*Allgemeine Gleichbehandlungsgesetz*) of 14 August 2006.

³⁶ See, e.g. the framework of electronic communications and energy.

achieving that aim are appropriate and necessary”³⁷. With respect to discrimination based on nationality, the importance of direct and indirect nature of unequal treatment is more nuanced. While the distinction seems to play a role in the free movement law³⁸, it is not directly mentioned in Directive 2006/123/EC on services in the internal market. Specifically, Article 20(2) of the Services Directive allows for differences in the general conditions of access to services based upon the service recipient’s nationality or place of residence – both directly and not – to be “directly justified by objective criteria”. This broad and ambiguous wording has resulted in a limited effectiveness of the said provision³⁹, which might explain why it is being generally overlooked in the anti-discrimination literature. To remedy this situation, the recently adopted Regulation 2018/302 on addressing unjustified geo-blocking and other forms of discrimination specifies the situations in online transactions where there can be no justification for different treatment based on the customers’ nationality, place of residence or place of establishment. The regulation, which will apply as of December 2018, thus appears to be introducing a *per se* ban on the specified manifestations of both direct and indirect discrimination in access to goods and services. In doing so, it takes an important step in recognizing and operationalizing the principle of non-discrimination in access to goods and services. Unfortunately – and not unlike other European building blocks of non-discrimination in consumer markets – the regulation does not appear to be AI-proof.

Indeed, as observed by Hacker, even though the horizontal dimension of the EU anti-discrimination law is comparably broad, its provisions on “access to goods and services” beyond specific sectors are of limited help to the subjects of discriminatory decisions reached by algorithms. This is particularly because of the “availability to the public” criterion, which appears to be difficult to reconcile with the growing degrees of personalization of both commercial messages about products or services and the products and services themselves. What is more, according to the author, considerations of algorithmic accuracy provide an easy justification that businesses can put forward to contest allegations of indirect discrimination. The difficulties faced by potential victims in holding the owners of algorithms accountable are further exacerbated by considerable enforcement deficits such as the lack of access to data and algorithmic models (Hacker, 2018). Similar observations have been made in the US literature with respect to the relevant anti-discrimination laws applicable to private law relationships, including the rather limited provisions concerned with consumer markets, such as the Equal Credit Opportunity Act of 1974 (Schmitz, 2017).

Overall, whereas the EU is criticized for suffering from a ‘justice deficit’ (Kochenov et al., 2015), it managed quite successfully to make the non-discrimination principle a genuine European value. At the same time, and even though the more recent developments in Europe have been more dynamic than in the US (De Búrca, 2012; Yoshino, 2011), there is also no unanimity in the EU as to how deep and in what ways the principle of non-discrimination should reach. The major field of application is still to be found in labor relations and ever stronger in the public domain outside and beyond contractual labor relations. The EU legislator did not manage to fully extend the non-discrimination principle to the sphere of consumption. Consumers might suffer from discrimination when the access to the market is barred or when the supplier selects the consumer for economic transactions on discriminatory grounds. At the same time, the building blocks of the existing European framework can be criticized for being a step, if not a mile, behind the evolving socio-technological reality (Hacker, 2018). Controversially, the most obvious discrimination in consumer markets is not social in nature but economic. D. Caplovitz coined the saying of the ‘iron law of capitalism’ ‘the poor pays more’ (Caplovitz, 1967). AI enables even more

³⁷ Directive 2000/43/EC implementing the principle of equal treatment between persons irrespective of racial or ethnic origin, Article 2(2)(b); Directive 2004/113/EC implementing the principle of equal treatment between men and women in the access to and supply of goods and services, Article 2(b); see also case C-83/14 *CHEZ*.

³⁸ See, e.g. case C-379/87 *Groener*, also: Reich, 2014, 66. Note, however, that the framework of non-discrimination on the grounds of nationality, with respect to free movement of persons, is specified by Directive 2004/38/EC on the right of citizens of the Union and their family members to move and reside freely within the territory of the Member States.

³⁹ SWD(2016) 173 final.

sophisticated forms of personalized advertisements and offers, which invites a renewed focus and perhaps a wholly new perspective on the principle of non-discrimination in private law relationships.

1.3.4. Consumer privacy

For decades data privacy and consumer policy stood side by side. Data protection lawyers aimed at protecting the privacy of the citizens. The judgment of the German Constitutional Court on the right to informational self-determination had an impact far beyond German and Europe (see generally Simitis, 2010). It paved the way for the development of a European policy of data protection in the 1990s and now for the adoption of the General Data Protection Regulation. The GDPR's applicability as of the 25th May 2018 is accompanied by a high degree of public awareness on the behavior of the old and the digital economy. The regulation, already mentioned before, defines ground rules not only for the relationship of the citizen and the state, but also of the citizen (and consumer) towards the supplier. As a result, data protection and consumer protection are coming closer. Consumer data protection turned into a new field of legal research, which oversteps the boundaries of two formerly separated fields (Helberger et al., 2017). There is a strong connection between traders' power over consumers and their ability to make sense of vast amount of data using algorithms.

The GDPR put in place quite a complex governance system, to ensure that personal data is being processed in a fair and lawful manner. We analyze its requirements in detail elsewhere⁴⁰, for here it suffices to state that data subjects should always know for what purpose and by whom what information about them is being used. This horizontal regulation does not, per se, stipulate what purposes data should not be used for – hence the other substantive principles and legal instruments will most definitely exist in a symbiotic relationship with the GDPR.

Consumer data protection impacts and affects not only the European consumers but also the consumers outside Europe. 'The Brussels effect' (Bradford, 2012) can be easily studied in the reaction of the American based big internet companies. They have to adjust their data protection policy to the European standards. It remains to be seen whether these companies will pursue a double standard (one for the Europeans, one for the rest of the world) and whether and to what extent the big players will monetarize data protection by offering different prices for services with and without data privacy.

⁴⁰ See: *CLAUDETTE meets GDPR Automating the Evaluation of Privacy Policies using Artificial Intelligence* (2018), available at: http://www.beuc.eu/publications/beuc-x-2018-066_claudette_meets_gdpr_report.pdf

2. Sector analysis: What does business use AI for, what are the gains and the risks?

To ground the research in the actual market practice, we decided to begin by running a survey of the purposes for which the companies are using artificial intelligence and other algorithmic techniques. In order to do so, we analyzed more than 100 websites and online media entries. We searched for the AI corporate use in the following sectors: finance and insurance, information services, energy and “smart solutions”, retail, autonomous vehicles, healthcare and legal. For each set of empirical findings, we asked the following questions:

- 1. For what purposes is business using AI in the given sector?** As a first step in moving away from books, ideas in our heads and general statements about “the AI”, we need to establish in precisely what ways business is employing artificial intelligence in a given sector of the economy. This part is descriptive and analytical. From this analysis we move to the normative observations, by asking:
- 2. What is the gain for business?** The development of AI tools is obviously costly, and so in order to understand why companies are using artificial intelligence to automate/realize certain tasks, we need to understand what it is that they are trying to achieve. Is it direct or indirect benefit? Are the “traditional” business activities being automated, or are companies able to do new things thanks to AI? Is the gain actual or potential? Knowing all this is necessary to ensure wise (potential) regulation: if imposing certain duties on business leads to more costly/ more difficult AI development and deployment, what is it that the economy will lose? Some gains, like security or fraud prevention, might be beneficial for society as a whole. Others, like automation of labor in call centers, might be solely in the interest of business. In short: what is business optimizing for? If we regulate, what could they lose?
- 3. What is the (potential) gain for consumers?** Clearly, consumers can also benefit from business’ deployment of AI. Even though the main purpose of this project is to identify *risks*, one should not forget that apart from risks there might be added consumer value; and combatting the risks might, as a side effect, reduce consumer benefits as well. These benefits might come from lower prices of goods and services (e.g. costs saved by the smart grid might translate into lower consumer prices), higher quality of goods or services (AI-powered financial products might be safer, or lead to higher revenues) or reduced search/transaction costs (with personalized offers/advertising; personal assistants automating purchasing decisions). In short: what are consumers gaining, directly or indirectly, from business using AI? What could they lose if this was constrained?
- 4. What are the risks to consumers?** This question is the core of this project. When business is clearly getting something, and when consumers *might* be benefiting as well, what is the cost that consumers pay for that? To answer this question, we analyze the concrete uses in concrete sectors, as opposed to a general armchair-scholarship simply listing Big Words like “autonomy”, “privacy”, “discrimination”. This should lead us to a specific map of risks, which in turn, analyzed jointly with the benefits, can allow us to pose concrete research questions to be addressed in Phase Two (chapter 4 of this report). Are the risks potential (maybe, one day) or actual and happening now? Are they predominantly individual in structure, or rather collective? Do they occur, given the technological inseparability, in a mode which poses particular new types of problems to the regulators and enforcers? In short: if regulation is a response to a concrete problem, what is the problem?
- 5. What are third party effects?** One should not forget that even though the primary focus of our analysis is the triad of relations: business-consumer-civil society; economic relations occur in a wider social, economic, cultural, political and currently essentially global context. There are externalities, some more, and some less straightforward. What are the impacts of business’ use of AI on labor situation, environment, access to knowledge, social values? In short: when using AI in specific cases, what are the third party effects of its mass deployment?
- 6. Any other reflections**

The results are presented below, sector by sector, in the form of a table for each field of business. We first answer the questions about each sector, and then list the sources we based our analysis on, as well as further interesting readings.

2.1. Finance and insurance

The financial sector, including the insurance sector, has been one of the first to attract public and scholarly attention when AI and other algorithmic techniques were introduced there. This was notably due to high-frequency trading and its potential role in the 2008 financial crisis, as well as credit scoring and potential for price discrimination. Financial institutions use artificial intelligence for pattern recognition, both in 1) data regarding security of their systems, 2) data about the markets (for more efficient algorithmic trading), and 3) data about consumers, for the creation of personalized financial products, risk assessment, and personalized pricing. In addition, financial institutions, just as many other sectors, employ chatbots for communication with consumers and consumer support, as well as making the personalized advice sound more “human”.

Sector	Finance and insurance
What is business using the AI for?	<ol style="list-style-type: none"> 1. Fraud prevention, e.g. with respect to insurance fraud – recognising patterns from a vast range of data sources to identify bad behaviour 2. Cybersecurity – analogous application 3. Customer support – chatbots communicating with consumers; systems allowing companies to prioritize customer messages (e.g. by classifying them according to the emergency and intensity of emotions) 4. Personalized and interactive financial advice – recognising patterns from a vast range of data sources to identify individual needs and risk profiles of consumers; the use of robo-advisors 5. Algorithmic trading / high frequency trading – the use of AI to predict the most optimal trading environment and the use of algorithms to automatically place orders on the stock exchange 6. Personalized insurance rates – assessment of risk based on collected data (e.g. collected by a mobile application, e.g. Snapshot introduced by Progressive to provide personalized car insurance)
What is the gain for business?	<p>Ad 1 and 2 – Better fraud prevention/security</p> <p>Ad 3 and 4 – Fewer employees needed for customer support and advice; improving reaction times, better prioritization</p> <p>Ad 4 – Establishing long-term relationships with the customers: the longer the customers remain with the company, the more data the latter has on them and (potentially) the better advice it is able to provide</p> <p>Ad 4 and 5 – Encouraging customers to become more active on the financial market (hence more profit), including in the area of stock trading</p> <p>Ad 6 – Improved risk assessment</p>
What is the gain for	<p>Ad 1 and 2 – Better protection against fraud/cyber attacks</p> <p>Ad 3 and 4 – Easier and faster communication with the company, available 24/7</p>

<p>consumers?</p>	<p>Ad 4 – Better access to financial products specifically tailored to one’s needs</p> <p>Ad 4 and 5 – Potentially more effective wealth management</p> <p>Ad 6 – Potentially lower insurance rates</p>
<p>What is the loss for consumers?</p>	<p>Ad 1 and 2 – using AI to prevent wrongful behaviour may not be controversial in some areas (e.g. fighting cyberattacks), but where the assessment is less clear and the system is also self-enforcing, due process arguments can be raised</p> <p>Ad 3 and 4 – Potentially lower quality of information provided; potentially more difficult to communicate with (human) employees</p> <p>Ad 4 and 5 – Responsibility for erroneous advice or for the wrongly placed orders unclear, risk of abuse / biased advice (robo-advisor maximising gains for the company, not for the customers), growing information asymmetry: companies can now gain better insights from consumer data, but consumers are not necessarily able to evaluate the quality of advice provided to them based on these new insights; how a particular conclusion was reached by the AI can even be unknown to the company itself (at the current stage of development AI systems based on neural networks are functioning as black boxes)</p> <p>Ad 6 – Privacy, potentially higher insurance rates (depending on the case)</p>
<p>Third party effects/ externalities</p>	<p>Potential job losses (customer support, financial advisors, analysts)</p> <p>New sources of data about consumers, potentially available to the government, law enforcement, hackers and blackmailers; potential risk to the stability of the financial system</p>
<p>Other</p>	<p>The use of AI to provide personalized services as well as automatic advice and support (chat bots) and to counteract bad behaviour is an overarching theme. Some further examples are provided below (see e.g. how AI is transforming advertising and retail, how it can be used to fight fake news or to facilitate IP enforcement).</p> <p>Other examples of personalization:</p> <p>http://victor-charles.com/web/en/page-home-en/</p> <p>Other examples of chatbots used for customer support:</p> <p>https://www.ibm.com/blogs/watson/2017/06/how-autodesk-sped-up-customer-service-times-with-watson/</p> <p>http://www.natural-talk.com/</p> <p>Aside from consumer markets, AI – possibly in combination with the blockchain – could be used for example to prevent fraud in different sectors of economy, e.g. by verifying the authenticity and origin of high value goods such as diamonds:</p> <p>https://www.ibm.com/blogs/research/2017/05/power-blockchain-watson/</p> <p>https://www.itnews.com.au/news/everledger-uses-blockchain-to-stamp-out-blood-diamonds-455537</p>

Sources/Further Readings:

- PwC, *Artificial intelligence in financial services*, “PwC”, December 2016, <https://www.pwc.com/us/en/financial-services/research-institute/artificial-intelligence.html>

- PwC, *Sink or swim: Why wealth management can't afford to miss the digital wave*, "Strategy&", 1 June 2016, <https://www.strategyand.pwc.com/reports/sink-or-swim>
- Accenture, *Banking Technology Vision 2017*, "Accenture", 2017, <https://www.accenture.com/us-en/insight-banking-technology-vision-2017>
- Banking Tech, *The future of banking is here: cognitive banking*, "FinTech Futures", 5 August 2016, <http://www.bankingtech.com/502582/the-future-of-banking-is-here-cognitive-banking/>
- Mariya Yao, *Chatbots Go Cha-Ching: The Looming Impact of A.I. In Finance*, "Forbes", 19 April 2017, <https://www.forbes.com/sites/mariyayao/2017/04/19/chatbots-go-cha-ching-the-impact-of-ai-in-finance/#4f7d778220e3>
- Q°emotion, *How a bank managed to increase customer satisfaction through emotions analytics*, "Q°emotion", <https://www.qemotion.com/how-to-use-emotional-analysis-for-banking-and-insurances-industries>
- BI Intelligence, *IBM to use AI to help banks with cybersecurity*, "Business Insider", 8 December 2016, <http://www.businessinsider.com/ibm-to-use-ai-to-help-banks-with-cybersecurity-2016-12?IR=T>
- Lecia Papadopoulou, *How Watson AI is helping companies stay ahead of hackers and cybersecurity attacks*, "IBM", 14 August 2017, <https://www.ibm.com/blogs/watson/2017/08/how-watson-ai-is-helping-companies-stay-ahead-of-cybersecurity-attacks/>
- Mariya Yao, *Can A.I. Defend Our Financial Institutions Against Hackers?*, "Topbots", 29th May 2017, <http://www.topbots.com/financial-services-institutions-artificial-intelligence-cybersecurity/>
- nanalyze, *6 Startups Using AI for Algorithmic Trading Strategies*, "nanalyze", 7 August 2016, <http://www.nanalyze.com/2016/08/artificial-intelligence-algorithmic-trading/>
- Olivia Solon, *World's largest hedge fund to replace managers with artificial intelligence*, "The Guardian", 22 December 2016, <https://www.theguardian.com/technology/2016/dec/22/bridgewater-associates-ai-artificial-intelligence-management>
- Gary Brackenridge, *Machine learning is transforming investment strategies for asset managers*, "CNBC", 6 June 2017, <https://www.cnbc.com/2017/06/06/machine-learning-transforms-investment-strategies-for-asset-managers.html>
- Michael Chui, James Manyika, and Mehdi Miremadi, *Where machines could replace humans—and where they can't (yet)*, "McKinsey Quarterly", June 2016, <http://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/where-machines-could-replace-humans-and-where-they-cant-yet>

2.2. Information services

"Information services" might be an imperfect term regarding its descriptive power, but in our analysis it has proven extremely efficient in "catching" new digital economy's business models of social media, online advertising, content hosting and screening, as well as facilitating communication. Information services use artificial intelligence and other algorithmic techniques for **advertising**. This might be *the* subject attracting most attention of consumer lawyers interested in artificial intelligence (however, with an overspill to public law, vide political advertising in the aftermath of the Cambridge Analytica scandal). Information services also use AI to offer more personalized and targeted business communication. In addition, AI is used for translations, content screening (to assess its quality, compliance with the platform's rules of conduct, compliance with intellectual property rules of copyright and trademark law), as well as powering communication via chatbots.

Sector

Information services

<p>What is business using the AI for?</p>	<ol style="list-style-type: none"> 1. Transforming digital advertising – recognizing patterns at the individual level and not based on profiling (segmentation); displaying ads related to actual preferences of the targeted consumer and not only to products he or she has viewed, his or her demographic data, etc. 2. Machine translation – In 2016 Google decided to carry out an overhaul of its translation service (Google Translate) and replace the code developed by engineers over a decade with a neural network (Google's Neural Machine Translation system); systems based on neural networks could potentially be applied in other areas where existing machine translation systems are used today (like in the case of Airbnb listings) 3. Fighting fake news – the use of AI to determine whether a trending topic on social media is factual or not; analysing and identifying fake audio and video 4. Automatic assessment of the quality of user-generated content (e.g. articles on Wikipedia) and, possibly, adding missing information (e.g. based on information available in different language versions) 5. Virtual assistants (text-based and voice-based chatbots), integrated, for example, into messaging apps (like Facebook Messenger)
<p>What is the gain for business?</p>	<p>Ad 1 – Advanced consumer targeting and ads personalization could allow companies to maximize returns on digital advertising.</p> <p>Ad 2 – More accurate translation tools could be used as stand-alone products (Google Translate) or be integrated into other products (commercialisation potential). By way of illustration, the use of AI by Google reduced translation errors by an average of 60%. This allowed Google to effectively compete with other IT companies such as the Chinese company Baidu.</p> <p>Ad 3 and 4 – Improving credibility and quality of online content</p> <p>Ad 5 – Attracting more users and gaining new insights about them, which can be monetized e.g. via advertising or product innovation</p>
<p>What is the gain for consumers?</p>	<p>Ad 1 – Exposure to more relevant content</p> <p>Ad 2 – More accurate machine translation, potentially substitutable for human translation. Machine translation is faster and cheaper than human translation.</p> <p>Ad 3 and 4 – Easier access to content of a potentially higher quality, lower risk of misinformation</p> <p>Ad 5 – Virtual assistants could quickly answer to various queries that consumers may have (e.g. what is the cheapest way to stream a particular TV show) thus further reducing their search costs</p>
<p>What is the loss for consumers?</p>	<p>Ad 1 – An even bigger impact on consumers' purchasing decisions</p> <p>Ad 2 – Machine translation actually “paid for” with personal data. More generally: growing dependence on a limited number of IT players who are able to afford major AI investments. If automatic translation is further integrated into other products, and virtually indistinguishable from human input, one may wonder what would happen if a translation is incorrect and based on this a consumer takes a purchase decision he would not have taken otherwise.</p>

	<p>Ad 3 and 4 – Pre-verification done by an algorithm could discourage consumers from evaluating the content (news) critically; furthermore, criteria used by the algorithms could lead to discrimination/censorship of certain content creators (e.g. new contributors on Wikipedia) or types of content (e.g. on politically controversial topics); algorithms could further stabilize past patterns, reproduce mistakes or be subject to abuse, all this leading to a distortion of the content available to consumers.</p> <p>Ad 5 – Companies not only gain access to, but also an understanding of consumers’ communications (with chatbot and with other users) and can use it, for example for advertising purposes.</p>
Third party effects/ externalities	<p>Potential job losses (e.g. translators)</p> <p>Impact on the freedom of expression</p> <p>New source of data about the consumers, potentially available to the government, law enforcement, hackers and blackmailers</p>
Other	<p>A related issue is the use of AI in the creative sector. Here, two main dimensions can be identified: AI as creator and AI as enforcer. As regards the former, it is observed that creative input can potentially also come from AI systems and not only humans. AI-generated works / inventions could be produced faster, with less effort and in large quantities. In the latter respect, AI could not only be used to recognize the patterns of illegal behaviour, but could also learn how the law (e.g. on copyright exceptions) is applied. This would make the IP enforcement not only more effective, but also potentially more balanced (although not necessarily more legitimate and transparent). See:</p> <p>http://www.osborneclarke.com/insights/defending-both-creators-and-the-public-will-ai-and-the-blockchain-transform-ip-management/</p> <p>https://script-ed.org/article/computers-as-inventors-legal-and-policy-implications-of-artificial-intelligence-on-patent-law/</p> <p>https://law.maastrichtuniversity.nl/ipkm/artificial-intelligence-ai-and-intellectual-property-ip-a-call-for-action/</p> <p>Overlap with data protection and IP law</p>

Sources/Further Readings

- Juniper Research, *AI machine learning to drive ‘real time bid’ advertising spend to \$42bn globally by 2021*, “Juniper research”, 5 September 2016, <https://www.juniperresearch.com/press/press-releases/ai-machine-learning-to-drive-%E2%80%99real-time-bid%E2%80%99-a>,
- See also: <http://www.campaignlive.co.uk/article/watch-m-c-saatchi-launches-artificially-intelligent-outdoor-campaign/1357413>
- John Mruz, *AI-to-AI communication in advertising increases brand safety and improves ad performance*, “Ibm”, 11 July 2017, <https://www.ibm.com/blogs/watson/2017/07/ai-to-ai-communication-helps-increase-brand-safety-and-drive-better-ad-performance/>
- Gideon Lewis-Kraus, *The Great A.I. Awakening*, “The New York Times”, 14 December 2016, <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html>, See also: original paper describing Google’s project <https://research.google.com/pubs/pub45610.html>

- Thomson Reuters, *Machine learning turns “fake news” on Twitter into old news*, “Thomson Reuters”, 27 March 2017, <https://blogs.thomsonreuters.com/answerson/machine-learning-fake-news-twitter/>, See also: <http://www.fakenewschallenge.org/>
- Rob Price, *CGI and AI are going to turbocharge 'fake news' and make it far harder to tell what's real*, “Business insider”, 28 July 2017, <http://www.businessinsider.com/cgi-ai-fake-news-videos-real-2017-7?IR=T>
- Aaron Halfaker, Dario Taraborelli, *Artificial intelligence service “ORES” gives Wikipedians X-ray specs to see through bad edits*, “Wikimedia”, 30 November 2015, <https://blog.wikimedia.org/2015/11/30/artificial-intelligence-x-ray-specs/>, <http://www.bbc.com/news/technology-34982570>
- Aaron Halfaker, *Investing in our shared future, supported by AI: Announcing the Scoring Platform team*, “Wikimedia”, 19 July 2017, <https://blog.wikimedia.org/2017/07/19/scoring-platform-team/>
- Alexei Oreskovic, *Facebook bought an AI startup that could turn its middling virtual assistant into a Siri killer*, “Business insider”, 31 July 2017, <http://www.businessinsider.com/facebook-acquires-ozlo-ai-startup-2017-7?IR=T>

2.3. Energy and “smart solutions”

Whether AI-powered systems can actually be described as “intelligent” is a fascinating matter for a debate with high importance for the public perception and politics, yet arguably of secondary importance for our analysis. What came to be commonly accepted though, at least on the level of language, is that AI-powered systems and products can be “smart”. Smartphones, smart homes, smart TVs, smart grids, smart cities – soon everything might start to be smart. Every now and then the public excitement rises when it is revealed yet another thing has been made “smart”, as for example with Amazon’s introduction of the “smart lock”, to allow its couriers to open consumers’ doors in order to safely (sic!) place the parcel inside a consumer’s house⁴¹. “Smartness” is not only a feature of consumer products though – whole grids, adjusting distribution of any media, ranging from energy to internet, can become (and in many places already have become) “smart”. Hence, in the “smart sector”, business uses AI for both creation of smart consumer products and services, as well as smart delivery of “traditional” supplies like energy.

Sector	Energy and smart solutions
What is business using the AI for?	<ol style="list-style-type: none"> 1. “Smart-grid” – real time analysis of consumer energy consumption patterns 2. Smart solutions for public spaces and commercial buildings (e.g. smart elevators: the use of AI to predict and suggest resolutions to potential problems) 3. Smart solutions for households (e.g. digital assistants responding to consumer requests such as Google Home, Amazon Echo, Apple’s HomePod or Baidu’s Little Fish as well as more autonomous home hubs which are currently being developed)
What is the gain for business?	<p>Ad 1 – More efficient use of the grid, less energy losses, lower costs of grid maintenance</p> <p>Ad 2 – Improved efficiency, fewer faults, easier maintenance</p> <p>Ad 3 – Multiple commercialisation opportunities (selling the device itself, partnering up</p>

⁴¹ <https://www.theverge.com/2017/10/25/16538834/amazon-key-in-home-delivery-unlock-door-prime-cloud-cam-smart-lock>

	with other companies and with application developers, gaining new insights about consumers and utilising them for advertising)
What is the gain for consumers?	<p>Ad 1 – Potentially cheaper energy</p> <p>Ad 2 – Depends on particular application (for elevators: less waiting time, fewer delays and malfunctions)</p> <p>Ad 3 – Facilitating the daily running (obtaining information, planning the day, making calls, playing music, shopping online), optimising the use of water, power, etc. (switching off the lights, adjusting the heating/cooling systems), improving safety (controlling locks and alarms), facilitating home maintenance. In the future: predicting problems before they happen and responding accordingly.</p>
What is the loss for consumers?	<p>Particularly Ad 1 and 3 – Business gains even more information about consumer’s habits (e.g. what amount of energy they are using), hence also when they are at home and when not, when they are sleeping, when they have guests etc. (privacy)</p> <p>Ad 3 – Risk of malfunctioning or hacking; problems with liability attribution</p>
Third party effects/ externalities	<ol style="list-style-type: none"> 1. Smart grid and smart homes could be good for the environment (more efficient use of energy means less pollution) 2. Potential job losses (e.g. maintenance) 3. New sources of data about the consumers, potentially available to the government, law enforcement (for example, in a US case authorities have asked Amazon to turn over data from a suspect’s Echo), hackers and blackmailers
Other	<p>Overlap with data protection law</p> <p>Possible combination of AI and blockchain</p>

Sources/Further Readings:

- Scott Carey, *How EDF wants AI to optimise its nuclear power stations and the smart home?*, “ComputerworldUK”, 11 May 2017, <http://www.computerworlduk.com/data/edf-eyes-ai-optimise-its-nuclear-power-stations-smart-home-3658843/>
- Constance Douris, *Balancing Smart Grid Data and Consumer Privacy*, “Lexington Institute”, 14 July 2017, <http://www.lexingtoninstitute.org/13750-2/>
- Jamie Condliffe, *What’s Next for AI Home Assistants*, “MIT Technology review”, 16 February 2017, <https://www.technologyreview.com/s/603672/whats-next-for-ai-home-assistants/>
- Michael Reilly, *Should an Amazon Echo Help Solve a Murder?*, “MIT Technology review”, 27 December 2016, <https://www.technologyreview.com/s/603278/should-an-amazon-echo-help-solve-a-murder/>
- Memoori Smart Buildings research, *The Innovative Startups that could bring AI & Blockchain to Smart Buildings*, “Memoori”, 2 March 2017, <https://www.memoori.com/innovative-startups-bring-ai-blockchain-smart-buildings/>, See also: <https://www.memoori.com/automated-smart-cognitive-past-future-smart-buildings/>

- Lynne Slowey, *Look who's talking: KONE makes elevator services truly intelligent with Watson IoT*, "IBM", 16 February 2017, <https://www.ibm.com/blogs/internet-of-things/kone/>, <https://www.ibm.com/watson/stories/kone-with-watson.html>, <https://www.kone.us/smart>
- -, *Meet MyxyPod. A voice assistant and multiroom speaker*, 2016/2017, <https://myxyty.com/myxypod-smart-home-speaker/>
- Freddie Dawson, *The House That Learns: Bringing Artificial Intelligence Into The Home*, "Forbes", 24 May 2016, <https://www.forbes.com/sites/freddiedawson/2016/05/24/the-house-that-learns-bringing-artificial-intelligence-into-the-home/#54f27c7f3fa3>, <http://aiport.ai-build.com/>

2.4. Retail

According to a popular story, a father discovered that his daughter was pregnant from baby product samples sent by a supermarket chain, which in turn predicted so based on her shopping patterns in the light of vast amounts of all customers' data. Whether this really happened does not really matter. What matters is that this story, just as many others, is extremely plausible. Retail selling – the very core of the consumer society, when the act of consumer *purchasing* happens – has been using artificial intelligence in a variety of ways. Firstly, adjusting the way the products are displayed, physically or virtually. We all know how in supermarkets crisps tend to be placed next to beer, water on the opposite side of the shop from bread, sweets on the level of children's eyes next to the cash desk, when they are bored waiting in line. These easy tricks can be now refined to a level of sophistication unseen before, based on millions of data points about consumers' purchasing patterns. Secondly, the communication with the customers has become much more personalized, especially where online shopping is concerned, with the boundary between (personalized) retail offers and advertising becoming more blurred than ever before. Finally, with the rise of "algorithmic consumers", "cyber-butlers", i.e. digital personal assistants, the very idea that it is a human making a purchase decision, might to a certain extent soon become outdated – with all the ensuing gains and losses.

Sector	Retail
What is business using the AI for?	<ol style="list-style-type: none"> 1. Providing insights for improved omni-channel retail solutions, i.e. solutions that combine brick-and-mortar stores and traditional delivery options with a variety of digital channels such as: in-store digital tools (interactive catalogues, intelligent mirrors, beacons, video cameras, robo-advisors), websites, mobile applications, customer cards, social media, chats bots, etc. 2. Creating personalized offers (e.g. coupons to be downloaded) and recommendations (e.g. which size or type of a dress should a consumer buy) based not only on previous purchases, but on wider and more precise data sets; providing real-time information about products which could be of interests to consumers (e.g. sending a mobile notification about a discount on umbrellas just as a consumer is approaching a retail shop on a day when rain is expected) 3. Understanding current trends (e.g. fashion trends based on Instagram photos) and creating corresponding products 4. Some digital assistants for households (like Amazon Echo/Alexa) can already be asked to place orders for products and services
What is the gain for	Improved merchandising operations, better management and monitoring of the supply chains, optimised inventory levels

business?	
What is the gain for consumers?	An easier and more flexible shopping experience
What is the loss for consumers?	<ol style="list-style-type: none"> 1. Retailers gain access to and understanding of a vast amount of data about particular consumers - also in brick-and-mortar stores (e.g. through beacons or video analytics allowing companies to tell which products people are looking at, and even where they are looking on the product) - and are able to exploit their weaknesses 2. Personalized offers could be discriminatory 3. Privacy issues 4. Risk of haphazard (uninformed) or erroneous purchases (e.g. as a result of a machine error or compliance with a child's request)
Other	Overlap with data protection law

Sources/Further Readings:

- Roger Bales, *How cognitive computing will revolutionize the retail industry*, “Ibm”, 15 May 2017, <https://www.ibm.com/blogs/systems/how-cognitive-computing-will-revolutionize-the-retail-industry/>, <https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=GBE03769USEN>
- ELEKS, *How Machine Learning Can Drive Retail Sales*, “NewCo Shift”, 31 May 2017, <https://shift.newco.co/how-to-use-machine-learning-to-sell-better-8f909ab425b4>
- Coherent Path, *Is ‘artificial intelligence’ the new ‘omnichannel’?*, “Coherent path”, 13 April 2017, <http://coherentpath.com/is-artificial-intelligence-the-new-omnichannel/>
- Vijayakumar Kabbin, *4 industries that will be transformed by machine learning in 2017*, “Information age”, 6 April 2017, <http://www.information-age.com/uk-public-trial-driverless-shuttle-bus-123465530/>
- Softbank, n/d, *Pepper (robot): For business*, “Softbank robotics”, <https://www.ald.softbankrobotics.com/en/solutions/business>
- Will Knight, *Amazon Has Developed an AI Fashion Designer*, “MIT Technology review”, 24 August 2017, <https://www.technologyreview.com/s/608668/amazon-has-developed-an-ai-fashion-designer/>
- Sarah Perez, *Starbucks unveils a virtual assistant that takes your order via messaging or voice*, “Techcrunch”, 30 January 2017, <https://techcrunch.com/2017/01/30/starbucks-unveils-a-virtual-assistant-that-takes-your-order-via-messaging-or-voice/>
- Jess Anderson, *Tesco and Estimote, pioneering grocery stores of the future*, “Estimote”, 8 March 2017, <http://blog.estimote.com/post/158169494355/tesco-and-estimote-pioneering-grocery-stores-of>
- Cosmose, *Omnicoookie connects offline stores with online ads*, “Cosmose”, <https://cosmose.co/> -, Not AI, but potentially a new source of data for AI analysis
- Heuritech, *Heuritech launches an Ai solution for Fashion & Beauty*, “Heuritech”, <http://www2.heuritech.com/>

- Harriet Taylor, *Lowe's introduces LoweBot, a new autonomous in-store robot*, “CNBC”, 3 August 2016, <https://www.cnbc.com/2016/08/30/lowes-introduces-lowebot-a-new-autonomous-in-store-robot.html>
- Andrew Liptak, *Amazon's Alexa started ordering people dollhouses after hearing its name on TV*, “The Verge”, 7 January 2017, <https://www.theverge.com/2017/1/7/14200210/amazon-alexa-tech-news-anchor-order-dollhouse>

2.5. Autonomous vehicles

Probably nothing excites the public opinion as much as the subject of autonomous vehicles – drones and self-driving cars. The former appear in discussions ranging from automated delivery of goods, to spying, to warfare; and the issues of privacy and security polarize the discussions. The latter fueled and gave new relevance to the endless discussion of the “trolley” problem, the question of whom a car should sacrifice, if it cannot avoid a collision. Unsolved questions that were purely theoretical until now have gained salience. With huge promises of increasing safety and efficiency of transport, self-driving cars also tend to be seen as a source of danger, disrupting the way in which transport has functioned for over a century now. A significant number of studies has been devoted to autonomous vehicles, which is why we decided to not really concentrate on them. Still, leaving them out of the general picture would leave a big gap in our map, so we conducted a little study of this subject as well.

Sector	Autonomous vehicles
What is business using the AI for?	<ol style="list-style-type: none"> 1. The use of AI in the development of autonomous cars (early stage) – self-driving cars not following the instructions provided by a programmer or relying on external analytics (i.e. sending data to a server or the cloud for analysis), but rather relying on an algorithm that has taught itself to drive and is able to respond to the situations on the road in real time. 2. The use of AI in the development of autonomous drones – drones not requiring a human operator, able to recognize and analyse their surroundings, including people below them (facial recognition), learning and improving themselves in the process 3. Combining the two: creation of wheeled, autonomous drones capable of switching between flying and driving
What is the gain for business?	Commercialising autonomous vehicles in new markets or using them to improve operations in existing markets: <ol style="list-style-type: none"> 1. Passenger transport: autonomous cars for personal use as well as autonomous taxis (for both urban and long-distance transport) 2. Aerial photography, carriage of goods: autonomous trucks and drones (in addition to numerous non-consumer uses such as industrial maintenance and servicing, law enforcement and military)
What is the gain for consumers?	Potentially faster, safer and easier transportation
What is the loss for consumers?	Safety and accountability:

	<ol style="list-style-type: none"> 1. Reliance on neural networks makes it very difficult to tell how a given result was reached – problematic in case of an accident or an error 2. Risk of hacking
Third party effects/ externalities	<p>Privacy of third parties (e.g. photographed by a drone)</p> <p>Damage caused to third parties (e.g. an autonomous car may prioritize the life, health and property of the vehicle’s owner or user over that of a passer-by)</p>
Other	Ethical issues: whom to sacrifice if necessary

Sources/Further Readings:

- Mariusz Bojarski, Ben Firner, Beat Flepp, Larry Jackel, Urs Muller and Karol Zieba, *By End-to-End Deep Learning for Self-Driving Cars*, “Nvidia”, 17 August 2016,
- <https://devblogs.nvidia.com/parallelforall/deep-learning-self-driving-cars/>, <http://www.nvidia.com/object/drive-automotive-technology.html>, <http://www.nvidia.com/object/uavs-drones-technology.html>
- Lex Davies, *Ford says it’ll have a fleet of fully autonomous cars in just 5 years*, “Wired”, 16 August 2016, <https://www.wired.com/2016/08/ford-autonomous-vehicles-2021/>
- Katie Collins, *Driverless taxis could be the bullet train of the Tokyo 2020 Olympics*, “CNET”, 22 August 2016, <https://www.cnet.com/news/driverless-taxis-could-be-the-bullet-train-of-the-tokyo-2020-olympics/>
- Johana Bhuiyan, *Uber wants to demonstrate its network of flying cars by 2020*, “Recode”, 25 April 2017, <https://www.recode.net/2017/4/25/15422592/uber-flying-cars-vtol-network-2020-dubai-texas>
- Simson Garfinkel, *Hackers Are the Real Obstacle for Self-Driving Vehicles*, “MIT Technology review”, 22 August 2017, <https://www.technologyreview.com/s/608618/hackers-are-the-real-obstacle-for-self-driving-vehicles/>
- Will Knight, *The Dark Secret at the Heart of AI*, “MIT Technology review”, 11 April 2017, <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/>, - discusses current research to make AI more explainable
- Tom Simonite, *AI-Powered Drone Will Follow You Around and Take Pictures*, “MIT Technology review”, 29 March 2017, <https://www.technologyreview.com/s/604009/ai-powered-drone-will-follow-you-around-and-take-pictures/>
- James Ryan, *Managing traffic in the skies*, 7 June 2017, <https://blog.x.company/managing-traffic-in-the-skies-494e2d992358>
- The Economist, T.S., *Why Uber’s self-driving car killed a pedestrian*, “The Economist”, 29 May 2018, <https://www.economist.com/the-economist-explains/2018/05/29/why-ubers-self-driving-car-killed-a-pedestrian>

2.6. Healthcare

Healthcare, just like privacy and data protection, is a subject-matter where a typical continental lawyer would at this point stand up and shout: “health is a constitutionally protected good, not a product! People receiving healthcare are patients, not consumers! How dare you apply market logic on such a

fundamental good as human health!'. In a way, we cannot disagree with this point – health is one of the most important goods protected by the legal system, and a situation in which those who can afford treatment end up being healthier than those who cannot, is socially undesirable. There is, indeed, some ethical dissonance in placing health in the same category as selling shoes. However, the social and market practice do not seem to be bothered by these considerations, and health services, as well as development and sales of healthcare equipment, have become an enormous sector of the contemporary economy. Therefore, admitting that health should be given a special place in market considerations, we decided to analyze this sector in the same manner as any other. Indeed, the possible amount of applications of AI in healthcare is immense. From diagnostics to optimizing hospitals' work and the usage of autonomous vehicles in life-saving operations, artificial intelligence is roaming around our hospitals. Let us hope that compassion will remain genuine.

Sector	Healthcare
What is business using the AI for?	<ol style="list-style-type: none"> 1. Diagnostics – examples <ul style="list-style-type: none"> • AI systems can be taught to diagnose from a case description and a clinical image (IBM Watson system has, for instance already been applied to examine mammography and cardiac patient imaging studies as well as tumor biopsies) • Streams, a system developed by Deep Mind (Google), is able to review test results for serious issues, such as acute kidney injury, and send smartphone alerts to the doctor. The system does not currently use machine learning although it is envisaged that in the future AI-driven alerts could also be delivered. Furthermore, AI could help predict future illnesses even before they manifest. 2. E-health solutions (mixed solutions, such as Babylon, in which medical advice is based both on AI analysis and on doctor's assessment; fully automated systems for purposes of self-diagnosis, health assistance (automatic nurses), medication management and even "therapy", e.g. Woebot (mental health chat bot), Quitxt (helps users quit smoking)) 3. Drone delivery of lifesaving medicines
What is the gain for business?	Commercialising AI systems in the healthcare sector Increasing efficiency of internal procedures Gaining access to even more consumer data, including sensitive data
What is the gain for consumers?	Ad 1 – More accurate analyses, potentially faster treatment Ad 2 – Cheaper and more accessible medical advice / "therapy" Ad 3 – Drone delivery is faster and doesn't rely on physical infrastructure
What is the loss for consumers?	Ad 1 and 2 – Risk of erroneous diagnoses and predictions, for reasons which might not be readily explainable (AI based on neural networks as a black box), less personal contact with medical professionals Ad 2 – Privacy, particularly where the system is developed by a commercial entity (e.g. Woebot belongs to Facebook)

Third party effects/ externalities	Risk of growing inequalities if only rich can afford AI-supported treatment
Other	Ethical issues: what to do in case of conflict between a human doctor and robotic judgment? Overlap with data protection law

Sources/Further Readings:

- Recode, *Full transcript: Zipline CEO Keller Rinaudo talks life-saving drones on Too Embarrassed to Ask*, “Recode”, 10 March 2017, <https://www.recode.net/2017/3/10/14875324/transcript-zipline-founder-keller-rinaudo-delivery-drones-too-embarrassed-to-ask>, <http://flyzipline.com/>
- Imaging Technology News, *Examples of Artificial Intelligence in Medical Imaging Diagnostics*, 16 December 2016, <https://www.itnonline.com/videos/examples-artificial-intelligence-medical-imaging-diagnostics>, <https://www.itnonline.com/videos/development-artificial-intelligence-aid-radiology/5239024570001>
- IBM, n/d, *Now patients can get personalized therapeutic options nationwide*, “IBM”, <https://www.ibm.com/watson/stories/quest-with-watson.html>,
- See also: <https://www.ibm.com/watson/health/oncology-and-genomics/oncology/>
- PubMed Health, *Behind The Headlines: Doctors 'vastly outperform' symptom checker apps*, “ITN”, 11 October 2016, <https://www.itnonline.com/videos/examples-artificial-intelligence-medical-imaging-diagnostics>
- Medical Futurist, n/d, *Artificial Intelligence Will Redesign Healthcare*, “The medical futurist institute”, <http://medicalfuturist.com/artificial-intelligence-will-redesign-healthcare/>
- Babylon, n/d, *Hand-picked doctors, supported by cutting edge technology* “Babylon”, <https://www.babylonhealth.com/>
- AiCure, n/d, *Artificial Intelligence for Continuous Patient Monitoring*, “AiCure”, <https://aicure.com/>
- Lora Kolodny, *Virtual nurse app Sense.ly raises \$8 million from investors including the Mayo Clinic*, 14 February 2017, <https://techcrunch.com/2017/02/14/virtual-nurse-app-sense-ly-raises-8-million-from-investors-including-the-mayo-clinic/>, See also: <http://sensely.com/>
- Mustafa Suleyman, Dominic King, *Independent Reviewers release first annual report on DeepMind Health*, “Deep mind”, 5 July 2017, <https://deepmind.com/blog/independent-reviewers-annual-report-2017/>
- Ciarán Mc Mahon, *What the Research Really Suggests About That Facebook Chatbot Therapist*, “Slate”, 18 July 2017,
- http://www.slate.com/blogs/future_tense/2017/07/18/what_research_says_about_woebot_the_facebook_chatbot_therapist.html, <https://www.woebot.io/>
- Paul Boutin, *The 7 Best Articles on Bots for Healthcare*, “Chatbots magazine”, 6 May 2017, <https://chatbotsmagazine.com/the-7-best-articles-on-bots-for-healthcare-28abc907528e>, <https://quitxt.org/>
- Recode, *Full transcript: Zipline CEO Keller Rinaudo talks life-saving drones on Too Embarrassed to Ask*, “Recode”, 10 March 2017, <https://www.recode.net/2017/3/10/14875324/transcript-zipline-founder-keller-rinaudo-delivery-drones-too-embarrassed-to-ask>, <http://flyzipline.com/>

2.7. Legal

AI-powered applications in the sector of legal services are booming. According to the Stanford CodeX Legal Tech index⁴², there are currently more than 1000 legal startups using AI to either develop in-house products or offer AI-powered products on the market. It turns out that a significant amount of tasks traditionally performed by lawyers can be to a large extent automated, even if we are still very far from creating a “robo-judge” and fully automating legal reasoning. Our own research confirms this claim in the sister-project called “CLAUDETTE” (<http://claudette.eui.eu>), where we tried to automate the analysis of consumer contracts and privacy policies of online platforms. These developments usually cause distress among young lawyers – “will we also lose our jobs?”. On the other hand, we believe that easier access to justice is to be welcomed, given how many people, despite legal professionals’ inflation, still cannot afford a lawyer, and how many tasks, especially in the third sector, do not get performed at all (e.g. abstract control of fairness in consumer contracts). However, just as in the case of any other AI-driven automation, there are risks we should be aware of.

Sector	Legal
What is business using the AI for?	<ol style="list-style-type: none"> Use by professional law firms: Systems such as Lex Machina or ROSS Intelligence are already able to retrieve relevant information from large sets of legal acts and case law and perform legal analytics. Further deployment of AI in this area could make it possible to form subjective interpretations and find out what a judge would likely do in the case at hand. Automated advice provided by chatbots such as DoNotPay
What is the gain for business?	Ad 1 – Reduces time and effort, allows lawyers to make better decisions as to which cases they want to take and adjust their fees accordingly
What is the gain for consumers?	<p>Ad 1 – Potentially more adequate fees (adjusted to the probability of success)</p> <p>Ad 2 – Access to basic (automatic) advice without a fee or at a significantly lower fee, for example a chat bot can successfully appeal parking tickets</p>
What is the loss for consumers?	<p>Ad 1 – Growing information asymmetry between lawyers and clients, potential for abuse, hindered access to legal advice for people identified as unlikely to win; lower chances of winning the case if one cannot afford the assistance of an AI-driven law firm</p> <p>Ad 2 – Legal advice provided by a bot might turn out to be of low quality</p>
Third party effects/ externalities	<p>Potential job losses (e.g. legal clerks, paralegals)</p> <p>The risk of stabilising existing patterns of legal argumentation and jurisprudence</p> <p>Access to justice</p>
Other	The impact largely depends on the extent of AI adoption. If AI systems were to become pervasive (like WestLaw or beck-online), then the problems of information asymmetry between lawyers and clients and access to justice could become a real issue. If, by contrast,

⁴² <https://techindex.law.stanford.edu>

AI systems were adopted by certain (more powerful) law firms only, other issues would arise (an even stronger position of a party which can afford the assistance of a law firm relying on AI, for example businesses in their disputes with consumers).

Sources/Further Readings:

- Artificial Lawyer, *The Lex Machina Story: From Start-Up to LexisNexis*, “Artificial Lawyer”, 21 March 2017, <https://www.artificiallawyer.com/2017/03/21/the-lex-machina-story-from-start-up-to-lexisnexis/>, <https://lexmachina.com/about/>
- Robert Ambrogi, *Lex Machina Adds Analytics For A New Area Of Law: Employment Litigation*, “LawSites”, 12 July 2017, <https://www.lawsitesblog.com/2017/07/lex-machina-adds-analytics-new-area-law-employment-litigation.html>
- Karen Turner, *Meet ‘Ross,’ the newly hired legal robot*, “The Washington Post”, 16 May 2016, https://www.washingtonpost.com/news/innovations/wp/2016/05/16/meet-ross-the-newly-hired-legal-robot/?utm_term=.f376504c6a06
- Danielle Muoio, *A 19-year-old created a free robot lawyer that has beaten 160,000 parking tickets*, “Business insider”, 28 June 2016, <http://www.businessinsider.com/joshua-browder-bot-overtakes-160000-parking-tickets-2016-6?IR=T>
- John Mannes, *DoNotPay launches 1,000 new bots to help you with your legal problems*, “Techcrunch”, 12 July 2017, <https://techcrunch.com/2017/07/12/donotpay-launches-1000-new-bots-to-help-you-with-your-legal-problems/>,
- <https://donotpay-search-master.herokuapp.com/>

3. Analysis by use: Data gathering and knowledge generation, automatic decisions, advertising, Personal Assistants

The effects of the sectoral analysis presented in the previous section demonstrate that, on the one hand, the gravity of the risks that consumer face differs from one part of the economy to another (the risk of being discriminated in advertising is not comparable to the risk of having a cancer misdiagnosed). On the other hand, there are certain common patterns one can observe. Having analyzed the empirical material, we have abstracted four general **uses** of AI by business:

1. data gathering and knowledge generation;
2. automated decision making;
3. advertising and other commercial practices;
4. digital personal assistants (or, more broadly, (ro)bots).

Important caveat: these “uses” do not constitute a clear-cut taxonomy. On the contrary, there might be quite some overlap in the analysis, given that these categories belong to different orders. For example: an AI-powered system displaying a targeted offer with a personalized price and form of presentation to a consumer, and reporting back on whether the consumer clicked, is simultaneously generating knowledge, taking an automated decision, and advertising. Hence, one could just as well imagine speaking about **perspectives** on the use of AI by business. We believe, however, that it makes sense to keep these analyses separate. The normative considerations that come to one’s mind when talking about decisions (fair/unfair) are different from those linked to commercial practices (fair/misleading/aggressive etc.). This said, we will not get defensive about this distinction – on the contrary, we would welcome any comments and feedback.

In this section we repeat the exercise from the previous one, i.e. study the way business uses AI and other algorithmic techniques, but instead of clustering uses in different sectors together, we cluster uses across sectors. One way to put it would be that as the previous section was based on vertical categories, this one is based on horizontal ones. Another way is to state that while the previous section was categorized by sectors, this one is categorized by uses.

How does this categorization fit with our claim that the intersection of consumer law and AI should be studied cases by case, bottom-up? **That’s where the big picture comes into play.** Our idea is that, to provide an extensive map of the ways in which businesses use AI, one could have a matrix with *sectors* on one axis, and *uses* on another. Clearly, the challenges stemming from data gathering and knowledge generation in healthcare or legal sector will be different from those in “smart” grids and houses, and might need different responses. However, it is possible that some uses can be studied across sectors (for example targeted advertising), and beg for special focus only in certain sectors (e.g. pharma). Hence, this approach allows us on the one hand to refine the analysis from the previous section, while on the other hand to notice some commonalities not visible after only the vertical study.

3.1. Data gathering and knowledge generation

Data gathering and knowledge generation might be the most fundamental “use” of AI by business, serving as a foundation for reflection and potential regulation. It was the availability of big data – the “dataquake” – that enabled high-paced development and mass deployment of AI in the first place; data not necessarily actively collected, but also generated as a side product of consumer use of online platforms and services (Alpaydin, 2016; Cristianini, 2016). Businesses use the data they gather to create profiles of consumers, use these profiles to display targeted content (ads/offers/content); and both to predict consumer behavior. In this sense, this first use is probably the most fundamental one, enabling the other ones.

Arguably, knowledge/information (a)symmetry has for long time been of interest to private and consumer lawyers (Akerlof, 1970). Classical economic models assumed perfect information. The

realization that information is not freely available to everyone equally (Coase, 1960) and – even more fundamentally – by nature local, dynamic and difficult to aggregate (Hayek, 1945) was one of the pillars on which the private law and regulation were based in the last half a century. Now, with big data and machine learning systems to make sense out of it, some of the most crucial assumptions are being challenged.

It was data that fueled the boom in recent developments in machine learning (Alpaydin, 2016). It is finding patterns in this data, as well as later using these patterns (profiles) for a variety of purposes, that seem to be the most fundamental aspect of the whole business AI enterprise. Data became a resource, valuable however only because machine learning systems are able to **turn it into knowledge**. The profession of data scientist is said to be the most promising job in the US right now. The availability of profiles/patterns, paired with advances in behavioral sciences, are a real game changer when it comes to knowledge generation right now.

All this gives rise to two general types of normative considerations. On the one hand, there are those that could be labeled “dignitarian”/“privacy for the sake of privacy”. These are based on the opinion that there are some things that companies should not know about consumers, regardless of what the effects of this knowledge might be. On the other hand, there are concerns that could be labeled “consequentialist” – concentrating on the ways in which this knowledge might be misused by companies, to the detriment of consumers. Both strands of the debate transpire through the existing scholarship, including some of the notable positions listed below.

A robust academic position regarding the use of artificial intelligence and other algorithmic techniques for profiling is the book edited by Hildebrandt and Gutwirth (2008): *Profiling the European citizen*. It provides an extensive overview of technical (what is profiling, how does it work), empirical (in what ways different stakeholders employ it) and legal knowledge (what legal goods are at stake). Within it, two chapters could serve as a marvelous source of knowledge for a researcher wishing to analyze the use of artificial intelligence by business from the point of view of consumer law. Firstly, the chapter by Kamp, Körffer and Meints, dealing with the questions of customer loyalty programs and scoring. Consumer partaking in loyalty programs, for example using loyalty cards in a supermarket, or collecting miles in air-traffic, become valuable sources of data for companies, both regarding the construction of profiles, and the later usage of these profiles for targeted commercial communications. Moreover, the customer behavior is being used as a source of knowledge to be deployed in risk assessment and fraud prevention in, among others, financial services. The authors analyze the problem first and foremost from the point of view of personal data protection law (at the time Directive 95/46/EC, now the GDPR); though their analysis can also serve as a valuable resource for others wishing to apply different normative standards to the issue at hand. The second position within this edited volume that could be of enormous value for a consumer lawyer studying business’ use of AI, is the chapter by Anrig, Browne and Gasson (2008). It provides a general overview of different types of algorithms used for the creation of profiles and finding patterns in behavior.

Several other interesting positions have been identified. While discussing the impact of data on digital markets, Kerber (2016) points out that the problems of privacy, competition and consumer protection cannot be separated, either analytically or for the purposes of regulation, and called for an integrated approach to scholarship, regulation and enforcement. The nature of digital markets is such that data will be used there to generate knowledge whether it is regulated or not. Since the existing legal framework might be insufficient, one could just as well consider the creation of new categories in this process. Costa-Cabral and Lynskey (2017) argue along similar lines and note that the values embedded in the data protection framework can be integrated into the competition policy and act as both internal and external constraint on competition law. Ezrachi and Stucke (2016) point to a threat of major online intermediaries (mega-platforms), which have access to vast sources of data, dominating the emerging digital markets, such as the on for personal assistants. Gal and Elkin-Koren (2017) attempt to allay these concerns and note the respective risk might be mitigated, among others, by the possible change of the “locus of data needed for the operations of algorithmic consumers from the Internet towards more

physical, and possibly less concentrated, loci (such as smart homes, smart cars, smart appliances, and smart clothes)” brought about by the Internet of Things. While this point is certainly valid, and experience shows that positions in the digital market can be much more fragile and elusive than they initially seem (Podszun and Kreifels, 2016), one can also argue that it is no longer data itself, but rather the ability to extract knowledge therefrom that allows key digital players to maintain their lead. This, in turn, requires access to advanced learning algorithms and sufficiently high computing power, which the new entrants may find difficult to afford.

On a different note, Barocas and Selbst (2016), discussing the big data’s “disparate impact”, point out that the existing data, from which knowledge is generated, contains historical biases. This is why this new “knowledge” (in inverted commas, since there is a chance for it not to be true) can contain and reproduce biases. Draper and Turow (2017), analyze in detail how data about consumer preferences can be used for audience construction. This data’s objectivity can be questioned not only because of the historic bias, but also as a result of its necessarily functional nature. An interesting comparative perspective has been offered by King and Forder (2016), who provide a survey of the approaches to knowledge building from the American and Australian perspective. Moving from knowledge generation by business for business to information provision, a very interesting analysis of the phenomena of filter bubbles and path dependence can be found in the piece of Desai (2015).

Further research into this area can go either in the more theoretical direction – i.e. what constitutes knowledge, what does it mean to *know* something in the algorithmic age – or in the direction of policy recommendations. The latter needs to be provided case-by-case, for each type of use separately. However, insights from other sectors can obviously inform and inspire further research projects.

Use type	Data Gathering and Knowledge Generation
In what way is business employing AI for this use?	<ol style="list-style-type: none"> 1. Structuring data 2. Finding patterns in data/ creating profiles 3. Quicker/more accurate response to the market needs 4. Enabling other uses
What is the gain for business?	<ul style="list-style-type: none"> - Ability to base the market/strategic decisions on much deeper and finer knowledge - Shorter reaction spans - Ability to engage in targeted commercial practices/ automate decision making process - Optimization of resources’ usage
What is the gain for consumers?	<ul style="list-style-type: none"> - Potentially, goods and services better fitting consumer needs - Potentially lower prices
What is the loss for consumers?	<ul style="list-style-type: none"> - Significant intrusion into privacy - Growing information asymmetry between consumers and business - Inability to understand business’ behaviour - Potentially, chilling effects on consumer behaviour

Third party effects/ externalities	<ul style="list-style-type: none"> - Higher market entry costs for businesses that do not have the knowledge/data - Natural data monopolies/ distortion of competition - Creation of new power centres (knowledge rich corporations) - Creation of new valuable data repositories, tempting for both governments and criminal groups
Other	Data has become a commodity and is de facto being treated like property, no matter how much private lawyers will insist it is not.

Sources/Further Readings:

- Nello Cristianini, *The road to artificial intelligence: A case of data over theory*, “New Scientist”, <https://www.newscientist.com/article/mg23230971-200-the-irresistible-rise-of-artificial-intelligence/>
- Louis Columbus, *Data Scientist Is the Best Job In America According Glassdoor's 2018 Rankings*, Forbes, January 29th 2018 <https://www.forbes.com/forbes/welcome/?toURL=https://www.forbes.com/sites/louiscolumbus/2018/01/29/data-scientist-is-the-best-job-in-america-according-glassdoors-2018-rankings/&refURL=https://www.google.it/&referrer=https://www.google.it/>
- Timothy Summers, *Facebook is killing democracy with its personality profiling data*, “The Conversation”, <http://theconversation.com/facebook-is-killing-democracy-with-its-personality-profiling-data-93611>
- Domingos, P. (2015). *The Master Algorithm*, Basic Books.

3.2. Automated decision making

While the concept of automated decision-making is commonly used in the literature, its meaning is not unambiguous. In a broad sense, automated decision-making can describe the very nature of IT-enabled algorithmic processes, which is producing outputs by means of executing a computer code (Article 29 Working Party, 2017a; Kroll et al., 2017). Admittedly, it is the fact that the underlying data collection and analysis as well as the subsequent procedural steps are performed automatically (by technological means) – and therefore more quickly and extensively than the same could ever be done by humans – that lies at the heart of the challenges investigated as part of this project. According to this understanding, algorithmic decision-making could thus refer to 1) automated data gathering and knowledge building and to 2) the performance of subsequent procedural steps – encoded in an algorithm or adjusted autonomously by artificial agents – with a view to reaching a predetermined goal. Obviously, such a perception gives rise to significant overlaps with other applications of AI in consumer markets investigated as part of this project. Once again, we would like to argue that this is not really a problem.

Indeed, one who embarks on the task of constructing a perfect definition of automated decision-making soon finds him or herself confronted with a range of difficult decisions to make. For a start, to what extent should one be concerned with a range of automated decisions leading to the production of a final output? Is it always possible to establish what the relevant final output is? And to what extent should one focus on whether there is a “human in the loop”? By way of illustration, in their recent report: *Public Scrutiny of Automated Decisions*, Upturn and Omidyar Network define automated decisions as “decisions made with the aid of systems that limit human judgment”. This strikingly anthropocentric notion follows the claim that “today’s automated decisions are not defined by algorithms alone. Rather, they emerge from automated systems that mix human judgment, conventional software, and statistical

models, all designed to serve human goals and purposes”. In other words, the wide empirical study conducted by both organizations suggests that it rarely happens that a decision would be taken solely by a machine. However, more recent legislative developments at the EU level, which directly engage with the topic of automated decision-making – particularly the General Data Protection Regulation 2016/679, focus on decisions taken *solely* by automatic means. From this it would seem that a full automation of a decision-making process creates specific risks with regard to the individuals concerned or, at least, makes the otherwise existing risks significantly more serious.

This brings us back to the two broader categories of the challenges that the use of AI by businesses poses to consumers. Just as with profiling and knowledge generation, on the one hand there is the “dignitarian” approach. According to this approach, the mere fact that a decision about an individual is being taken solely (or largely) by a machine is normatively suspect and should be counteracted, regardless of what the substance of that decision is. On the other hand, there are more liberal positions, which would rather be opposed to potentially negative consequences of such automated decisions. The former is linked to the broader effects of automatic decision-making upon personal autonomy and dignity (addressed, among others, by Gal and Elkin-Koren, 2017; Brownsword, 2017; André at al., 2018 and Mik, 2016). The latter, by contrast, is more concerned with the possibility of erroneous, discriminatory or exploitative outcomes (Brauneis and Goodman, 2018; Bar-Gill, 2018; Kim, 2017a; Barocas and Selbst, 2016; Zarsky, 2016; Vladeck, 2015; Citron and Pasquale, 2014; Schmitz, 2014). None of these approaches should be treated as prevailing globally.

There are clearly different types of automated decisions that might require different normative responses, depending on the matter to be automatically decided, and the legal good at stake. Some decisions might be those where we would like to have a human in the loop given how **grave** a matter at hand is. When it comes to sentencing, visa denial or medical treatment, we might think human dignity requires that an individual has to be there to decide about the fate of an individual, for purely deontological reasons. Naturally, consequences of unjust decisions related to matters of especially high legal or socio-economic gravity might be also especially grave. Not surprisingly therefore, this is also the area upon which the scholarship tends to focus (on the use of algorithmic techniques by employers and by public authorities see e.g. Kullmann, 2018; Kim, 2017b; Brauneis and Goodman, 2018; Dressel and Farid, 2018; Barrett, 2017; Rich, 2016). In other context, like online ads or amateur machine translation, we might be fine with certain decisions being taken just by machines. However, even when **gravity is high**, there are some spheres where a human cannot be kept in the loop for technical reasons (**too quick to consult**), like the decision of self-driving cars on whom to sacrifice (not to mention autonomous weapons – fortunately outside the realm of consumer law, for now). Then, when it comes to advertising or commercial practices, there might be actions **too frequent to consult**, as billions of decisions are taken daily. As always, the normative assessment should be conducted on a case-by-case basis.

Interestingly, the logic behind specific rules on automated decision-making found in the General Data Protection Regulation reflects both of the aforementioned perspectives, even if the instrument is generally recognized as imperfect. Indeed, Article 13(2)(f) requires the data controller to provide the data subject with the information on the existence of automated decision-making, including profiling, at the time personal data is obtained (yet, notably, not at the time when the decision is reached). The data controller should also provide the data subject with meaningful information about the logic involved, the significance and the envisaged consequences of such processing for the data subject. Furthermore, pursuant to Article 22(1), the data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or *similarly significantly* affects him or her. This right – already constrained by the “gravity” factor – is subject to further limitations laid down in the subsequent paragraph. These include situations where a decision is necessary for entering into or performance of a contract or is based on the data subject’s explicit consent. Even in that case, however, the GDPR obliges the data controller not to base automated decisions on special categories of personal data, such as data revealing racial or ethnic origin and

political opinions or health data (with limited exceptions), as well as to implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller. The right to be informed about automated decision-making, the right to object thereto as well as the right to obtain human intervention are all concerned with the deontological dimension, i.e. the automatic nature of the processing *per se*. Further rights assigned to the data subject – to express his or her point of view and to contest the decision – already hint at the consequential dimension, i.e. the specific decision reached. The GDPR in itself, however, does not provide a viable solution to this set of concerns, as observed by a number of commentators (Wachter et. al, 2017; Vedder and Naudts, 2017; cf. BEUC, 2018).

To sum up, we recognize that the question whether decisions are being taken with or without human involvement can be an important factor to bear in mind. However, what appears even more essential for the understanding of challenges that automated decision-making brings to consumer markets are the concrete purposes for which the investigated techniques are being used. Consequently, the remaining part of this section as well as the subsequent ones (in more specific contexts) will focus on the use of software agents, including AI-powered, that **automatically adjust business conduct in order to increase the likelihood of consumers purchasing a product or a service, or to otherwise optimize commercial gains of traders**. Automated decision-making in that sense often results from the previous processes of automated data gathering and knowledge building, addressed in section 3.1. above (similarly Article 29 Working Party, 2017a).

As regards the business use of automated decision-making, a distinction might be made between algorithmic decisions taken at a pre-contractual, contractual and post-contractual stage. The former encompass targeted and personalized advertisements and offers, including personalized prices (Bar-Gill, 2018; Steppe, 2017; Ezrachi and Stucke, 2016). Algorithmic decisions taken at this stage may also result in the refusal to serve a specific consumer – a decision of which the effects upon the consumer may be particularly significant (using the GDPR's terminology) when the lack of such access affects the consumer's ability to fulfill his or her essential social needs. The decisions taken at the contractual stage may refer, in particular, to personalized quality (e.g. of customer service; Schmitz, 2014; cf. Bar-Gill, 2018). This thread of the discussion is also linked to the so-called smart contracts, in which (in particular) contractual performance can become automated (Mik, 2017). However, both now and in the foreseeable future, smart contracts are not likely to be used in business-to-consumer relations to an extent higher than marginal. The same seems to be true for the use of algorithmic decision-making techniques at the post-contractual stage, particularly in dispute resolution (for a further discussion see Barnett and Treleaven, 2018; Larson, 2010). The focus of the present report thus remains on automated decisions reached before and during contractual relationships, other than smart contracts. It is also worth noting that in the long-term relationships, which appear to be increasingly prevalent in the digital economy (for smart products see: Helberger, 2016), elements associated with all contracting stages become interrelated. At the same time, boundaries between information and advertisement (or, as Helberger puts it, between “informing, nudging and outright manipulation”) become blurred (on a similar problem of drawing a line between health content and commercial content in the context of mHealth apps, see Sax et al., 2018). This is particularly true with respect to the personal digital assistants, which are discussed separately in section 3.4 below, following an additional section 3.3. on targeted and personalized advertising. As already explained, we chose to keep these discussions separate in order to underline the specific issues arising in these contexts, such as exerting technological influence and outsourcing decision-making powers.

A challenge arising from the business use of automated decision-making upon we wish to focus here is the use of software agents to **coordinate or personalize contractual conditions**, most notably the prices charged to particular consumers. The issues posed by pricing algorithms are linked to price coordination, on the one hand, and price discrimination, on the other.

The idea of algorithmic price coordination refers to the deployment, by different undertakings, of pricing algorithms, which continually monitor and adjust to each other's prices. The topic has met with a

particularly strong interest in the competition law literature. By way of illustration, Ezrachi and Stucke draw a number of scenarios – some of which they consider more problematic than others (Ezrachi and Stucke, 2017). The least controversial case is the one in which computers are used to directly execute human instructions (“the Computer as a Messenger”). At the other end of the scale are those situations in which the measures taken are determined autonomously by machine-learning algorithms and only the target – e.g. profit maximization – is being set by humans (“Digital Eye”). In this latter case a market outcome similar to a prohibited collusion among competitors cannot be ruled out, yet a breach of competition law – particularly the existence of an anticompetitive intent – might be difficult to demonstrate (Ezrachi and Stucke, 2017; Janka and Uhler, 2018). At the same time, there seems to be no uniform view in the literature as regards the extent and severity of this issue and, therefore, also as regards the relevant policy implications. While certain authors believe that there is a need to reconsider some of the embedded assumptions of competition law (Mehra, 2015), others regard the respective challenges as overstated and draw attention to the use of alternative policy tools, such as consumer protection laws (Lindsay and McCarthy, 2017).

Algorithmic price coordination should be distinguished from algorithmic decision-making processes resulting in price discrimination. In the economic theory the notion of price discrimination is associated with a practice of offering identical products or services to different consumers at different prices where such differences are not motivated by different cost structures, e.g. different costs of supply (Stigler, 1968; Levine, 2002). The term “discrimination” used in this context thus has a different meaning than in anti-discrimination law, as discussed in section 1.3.3 above. With the rise of big data, businesses are gradually gaining the capacity to adapt prices so that virtually every individual consumer can be charged the maximum price he or she is willing to pay. Such a phenomenon, referred to as first-degree price discrimination, is perceived particularly critically by the European data protection authorities (European Data Protection Supervisor, 2015; Article 29 Working Party, 2013). Analyses carried out from the economic point of view appear to be less critical and more nuanced (Levine, 2002; Fudenberg and Villas-Boas, 2012; Rayna et al., 2015; Thierry et al., 2015). In this respect a particularly interesting contribution has been provided by Oren Bar-Gill, who investigates, with the help of an economic analysis, in which market conditions price discrimination may result in both consumer harm and efficiency losses. The author explores the phenomenon of price discrimination with much nuance and proposes a range of respective policy options, such as increased privacy protection and enhanced enforcement of the rules against misleading and aggressive advertising, along with more unorthodox proposals like personalized price caps and personalized disclosures (Bar-Gill, 2018; see also Porat and Strahilevitz, 2014; Hacker, 2017).

Automatic decision-making leading to differences in prices charged to particular consumers may also go beyond price discrimination in the abovementioned understanding. This refers, in particular, to the markets where costs borne by the traders in connection to a transaction are heterogeneous, like in the case of credit or insurance. In this case, businesses have even higher incentives to invest in algorithms aimed to determine whether individual consumers qualify for a product or service and if so, on what terms the product or service should be offered to them. Even though such a differentiation may be considered less objectionable from the economic point of view and may even result in societal benefits (e.g. reduced human bias in consumer vetting), the risks posed by it for consumers cannot be disregarded. Indeed, many commentators point to the problems of biased datasets, erroneous or discriminatory categorizations as well as the risk of perpetuating inequalities and cycles of poverty (Schmitz, 2014; Chander, 2017; Feldman et al., 2015). The possibility of (mis)using insights about consumers to maximize business gains cannot be excluded either (see: Zarsky, 2016 on unfair transfer of wealth). The problem becomes even more significant if algorithms are based on machine learning and, according to the current levels of knowledge, essentially function as black boxes (Pasquale, 2015; Vladeck, 2016; Perel and Elkin-Koren, 2017; cf. Kroll et al. 2017, Rudin, 2014).

Policy responses to the concerns raised by algorithmic decision-making proposed in the literature are remarkably diverse. More often than not several complementary measures are recommended. The

specific proposals range from the calls for greater transparency (Schmitz, 2014), over technological solutions (Kroll et al., 2017; Doshi-Velez and Kortz, 2017), to *ex post* review (Kim, 2017a; Perel and Elkin-Koren, 2017) and social policy interventions (Zarsky, 2016 on government-based loans; Barocas and Selbst, 2016 on algorithmic decisions being “too accurate”, i.e. exposing societal implications of persisting inequalities).

Use type	Automated Decision Making
In what way is business employing AI for this use?	<ol style="list-style-type: none"> 1. Using AI-powered systems to partly replace humans decision makers 2. Using AI-power systems to take into account much bigger amount of data/knowledge 3. Automating the “mechanical” parts of the decision-making process, in order to leave the creative parts to humans
What is the gain for business?	<p>Ad 1 Increasing the pace of the decisions: ability to take business decisions much quicker</p> <p>Ad 2 Increasing the quality of the decisions: ability to base the decisions on a much higher amounts of data/ much more accurate knowledge;</p> <p>Ad 3 Potentially increasing objectivity of the decisions, removing undue influence, emotions, human limitations etc</p>
What is the gain for consumers?	<p>Ad 1 Potentially shorter waiting times</p> <p>Ad 2 Potentially more accurate decisions</p> <p>Ad 3 Potentially more objective/unbiased decision</p>
What is the loss for consumers?	<p>Risk of falling into historical bias</p> <p>Inability to question the automated systems decision</p> <p>Risk of unfair decisions, unduly favouring the business’ needs</p>
Third party effects/ externalities	The excluded/ historically discriminated against social groups risk deepening the exclusion
Other	There is a need to re-think the concepts of a decision, of agency

Discussion of selected papers

Bar-Gill, O. (2018). Algorithmic Price Discrimination: When Demand is a Function of Both Preferences and (Mis)perceptions. University of Chicago Law Review, Forthcoming

The paper by Bar-Gill investigates the issue of price discrimination in the age of big data. The analysis is based on the understanding of price discrimination, which does not, in itself, carry a negative connotation and simply refers to the differentiation of prices charged by a given supplier to particular consumers. The reasons for such a differentiation, in particular whether or not the differences in prices charged can be attributed to the different costs of supply, do not appear to be decisive for the definitional purposes of this paper. That said, the main part of the author’s analysis is, indeed, devoted to the markets where sellers offer homogeneous products to heterogeneous consumers and charge personalized prices that correspond to each consumer’s willingness to pay (WTP). One of the key insights of the paper is

the understanding of the consumer's WTP as a function of consumer preferences and (mis)perceptions. This finding has significant implications for the welfare effects of price discrimination, which are analyzed subsequently. According to the author, when algorithmic price discrimination targets preferences, it harms consumers but increases efficiency. On the other hand, when price discrimination targets misperceptions, and particularly demand-inflating misperceptions (which seems to be the most likely market scenario), not only the consumer welfare loss decreases, but so does the efficiency. According to Bar-Gill, in those circumstances legal intervention may be warranted. Among several policy options the author points to the increased privacy protections and enhanced enforcement of the rules against misleading and aggressive advertising, along with more unorthodox proposals like personalized price caps and personalized disclosures.

- **Problem addressed:** Big data provides traders with insights not only about consumers' preferences, but also about their misperceptions (cognitive biases etc.). The insights thus obtained can be used to establish individual consumers' willingness to pay and charge personalized prices so as to maximize traders' gains.
- **Key finding/argument:** Price discrimination, which relies on both consumer preferences and consumer misperceptions, not only decreases consumer surplus, but also reduces efficiency.
- **Solution proposed:** Policy responses could target either price discrimination or consumer misperceptions. The former include: direct prohibitions of price discrimination, facilitation of price arbitrage, improving transparency to trigger fairness-based consumer backlash, limiting sellers' access to information about consumers' willingness to pay (increasing privacy protections and data security) or setting personalized price caps. The latter include, in particular, more robust tools against unfair and deceptive practices and personalized disclosures.

Citron, D. K., & Pasquale, F. (2014). The Scored Society: Due Process For Automated Predictions. *Washington Law Review*, 89, 1.

The authors investigate the automated scoring systems, which have an increasing impact on crucial life opportunities of consumers/citizens. According to Citron and Pasquale such systems are not free from bias and can lead to arbitrary and discriminatory outcomes, as demonstrated by the case study of credit scoring. Nevertheless, methods and logic of predictive systems remains obscure (often guarded by trade secrets). The problem of obscurity becomes even more significant when machine-learning techniques are applied. The authors thus call for the integration of "cognitive perspective of AI", by which they appear to understand the introduction of human values such as fairness and justice to the picture. More specifically they refer to the concept of "technological due process" and require that both the scoring algorithms themselves ("technology-driven rulemaking") and the individual decisions taken on that basis ("technology-driven adjudication") are subject to review and revision. The values of due process to which they refer are those of transparency, accuracy, accountability, participation and fairness.

- **Problem addressed:** Automated scoring systems leading to arbitrary judgments and discrimination of historically subordinated groups. Broader social implications such as a widening gap between the privileged and disadvantaged (similarly, Schmitz, 2014). Although their influence is growing, algorithms cannot be meaningfully checked, among others for reasons of opacity.
- **Key finding/argument:** Need for "technological due process" – procedural framework to ensure that the scoring algorithms and subsequent decisions are subject to review and revision.
- **Solution proposed:** Individual rights available already at the stage of data gathering, access to data sets pertaining to specific individuals. Ideally public or, alternatively, confidential review of the logics of predictive scoring, including the source code (e.g. by trusted third parties). Informing scored individuals when their scores are communicated to subsequent decision-makers. Providing consumers with the tools for interactive modeling – ability to see the changes in one's score across the range of hypothetical situations). Allowing consumers to challenge adverse decisions affecting them. Licensing and auditing requirements for scoring systems at the stage when scores are used

in decision-making (at least in critical settings like employment, insurance, health care). Oversight by regulators, notably the Federal Trade Commission (under its authority to combat unfair trade practices) – access to scoring systems, testing on hypothetical examples by IT experts, issuing impact assessments evaluating the system’s negative effects and identifying risk mitigation measures.

- **Related literature:** Zarsky, 2014; Pasquale and Citron, 2014 (responses).

Kim, P. T. (2017a). Auditing Algorithms for Discrimination. University of Pennsylvania Law Review Online, 166(1), 189.

The author of the contribution responds to Kroll et al. (see below) by pointing to the limits of technical mechanisms for ensuring accountability of algorithms and arguing that auditing should remain an important tool to detect and counter discrimination. As a starting point, Kim observes that there is no consensus regarding the meaning of discrimination and that it may not be possible to satisfy the different definitions simultaneously. For example, the decisions made by a system that seeks to predict outcomes will differ depending on whether nondiscrimination is defined as 1) equalizing the proportion of correct negative and correct positive predictions for each group, or 2) equalizing the proportion of false positives or 3) false negatives across the groups. More importantly, the author notes that biased decisions may result not from the code or from implicit bias in the input data, but rather from the broader social processes. According to Kim, the use of randomization, advocated by Kroll et al., not only does little to remedy this situation, but can also be difficult to accept in many social contexts. By way of illustration, while incorporating randomness to establish which online ads are more likely to generate clicks may not be particularly problematic, the same cannot be said about the use of this technique to predict job performance, credit worthiness or recidivism (context matters). Kim furthermore takes issue with another strategy, discussed by Kroll et al. – the one drawn from the notion of “fairness through awareness” developed by Dwork et al. (2012). The author notes that it is not only exceedingly difficult to identify “all relevant features” which determine “how different people are” – also measuring them in an unbiased way might be a daunting task. Once again, the problem seems to be the one of societal origin and not likely to be remedied by technological tools. Consequently, Kim underlines the importance of ex post auditing of algorithmic decision-making processes. In her view, contrary to the suggestion of Kroll et al., the judgment of the U.S. Supreme Court in *Ricci vs. DeStefano* does not preclude ex post correction of detected discriminatory outcomes, especially the prospective ones. On the whole, although Kim observes that “designing algorithms to be nondiscriminatory is by far the preferable practice”, the author warns against an excessive reliance on these techniques.

- **Problem addressed:** Discriminatory outcomes of automated decision-making, which cannot be addressed by technological tools alone.
- **Key finding/argument:** The sources of certain discriminatory outcomes lie in the broader societal processes. Technological tools discussed by Kroll et al. are not suitable to address them. The judgment in *Ricci* does not preclude ex post auditing.
- **Solution proposed:** Ex post auditing should remain an important part of the framework for countering discrimination by algorithms.
- **Related literature:** Kroll et al., 2017.

Kroll, J. A., Barocas S., Felten E. W., Reidenberg J. R., Robinson D. G., & Yu H. (2017). Accountable Algorithms. University of Pennsylvania Law Review, 165, 633.

The article by Kroll et al. discusses the specific technical tools for improving the accountability of algorithmic decision-making. The authors underline the limits of transparency and ex post auditing and argue for an ex ante design of algorithms so as to ensure procedural regularity (due process) and assist in the achievement of substantive policy goals. Noteworthy mechanisms include techniques aimed to verify whether several investigated decisions were reached according to the same and previously

disclosed procedure (cryptographic commitments, zero-knowledge proofs) and whether random inputs were chosen fairly (incorporating seed values). The authors further observe that incorporating well-designed randomness can maximize the gains from learning from experience and hence reduce the negative effects of past prejudice and implicit bias reflected in the input data. Several techniques for furthering substantive fairness, defined in various ways, are also illustrated. One of such techniques relies on the theoretical model of “fairness through awareness”, developed by Dwork et al. (2012), which requires that similar individuals have a similar chance of receiving any possible outcome. Other approaches are based on the concepts of “fair affirmative action”, “fair representation”, “regularization”, “fair synthetic data” or “interpretability” of machine learning systems (on interpretability, or the role of explanation, see also Doshi-Velez and Kortz, 2017). The authors then briefly discuss the impact of applicable US antidiscrimination law and point to the tensions posed by algorithmic decision-making to its foundational concepts such as disparate treatment and disparate impact. A reference is also made to the judgment of the US Supreme Court in *Ricci v. DeStefano*, which is viewed as a possible legal obstacle to ex post corrections of discriminatory algorithms (note, however, that a different reading of *Ricci* has been advocated by Kim, 2017a – see above). The paper concludes with recommendations addressed at both computer scientists and policymakers.

- **Problem addressed:** A growing number of decisions are taken by technological means, yet accountability mechanisms have not kept up with technology. Transparency as well as ex post review and oversight are in themselves neither optimal, nor sufficient.
- **Key finding/argument:** Technological tools exist that can improve accountability of algorithms by making it possible to verify, ex post, whether automated decisions comply with the legal standards. Tools to ensure procedural regularity are already within reach. Similar approaches can also increase accountability for substantive policy goals like non-discrimination.
- **Solution proposed:** The discussed techniques should be implemented ex ante, at the design stage. Further research regarding specific tools as well as an enhanced cooperation between computer scientists and policymakers are necessary. Administrative agencies could become engaged in developing guidelines for software developers.
- **Related literature:** Kim, 2017a (response); Dwork et al., 2012.

Schmitz, A. J. (2014). Secret Consumer Scores And Segmentations: Separating “Haves” From “Have-Nots”. Michigan State Law Review, 1411.

In her article, Schmitz investigates the broader societal implications of consumer scoring. The discussion goes beyond non-discrimination on the basis of specific characteristics, but is concerned more generally with the very use of predictive segmentations and scores by companies “to assess each consumer’s likely value to the company and to decide what offers and remedies each consumer deserves in the company’s assessment”. The author argues that such a differentiation in price and quality can perpetuate the cycles of poverty and increase power imbalance between the powerful “haves” and the disempowered “have-nots”. The article goes on to discuss the extent to which the identified concerns are addressed by legal mechanisms available in the US (Fair Credit Reporting Act, Federal Trade Commission Act, Dodd-Frank Act establishing the Consumer Financial Protection Bureau and federal antidiscrimination law, including the Equal Credit Opportunity Act). The author does not consider this framework to be sufficient and argues for increased privacy protections, most notably related to transparency, opting-out from data processing, and dispute resolution. Attention is also drawn to the role of the FTC and a call is made for stricter auditing and accountability of data brokers.

Zarsky, T. (2014). Understanding Discrimination in the Scored Society. Washington Law Review, 89, 1375.

The article is a response to the paper by Danielle Citron and Frank Pasquale, *The Scored Society* (Citron and Pasquale, 2014; discussed before). The author concurs with their overall argument that the use of

algorithms for purposes of consumer scoring brings about significant societal challenges. What he is missing in the original contribution is a deeper scrutiny of the discrimination-based concerns associated with this process. To fill this gap, Zarsky offers an analysis of the notion of discrimination (in theory, not in law and doctrine), identifies both the opportunities and threats brought by consumer scoring and outlines the possible policy responses. Because this dimension of anti-discrimination law has not been addressed in much detail before, the subsequent discussion of Zarsky's paper is more extensive.

The article discusses the strength of specific fairness-related arguments, which underlie the discrimination-based concerns, across a variety of contexts. These arguments can be either deontological (discriminatory intent), or consequential (discriminatory outcome: social segregation and stigmatization). The most extensive discussion is devoted to discrimination against protected groups, in which case the rationales against discrimination appear to be most salient. In this respect the author distinguishes, yet again, between explicit discrimination, implicit discrimination and instances of disparate impact. The overall message seems to be that justifications for prohibiting discrimination are strongest in the former case and weakest in the latter, even though the boundaries between particular categories may not be easy to draw. One of the noteworthy observations is that – from the point of view of investigated rationales – implicit discrimination can often be equally unacceptable as the explicit one. Consequently, the author calls for an expansive notion of intent, which also includes certain reckless or negligent actions, such as the reliance on tainted tools and datasets or the use of blatant proxies.

The author goes on to discuss the specific consequential justifications, which can be invoked against discrimination via disparate impact (i.e. where no discriminatory intent can be found). He begins with the theory of social segregation, which he considers particularly relevant in the context of credit and employment as the areas in which “lower ratings and scores applied to a specific group quickly translate into transfers of wealth and changes in social structure”. According to the author, the argument does not work with equal strength for other contexts such as “advertising, marketing, and perhaps insurance”. Here, however, the stigma-based concerns may be more serious. The specific assessment depends on “how visible and salient the process is in the eyes of those discriminated against, as well as other segments of the public”. This, in turn, is linked to the transparency of algorithmic processes and the gravity of the social stigma attached to a given group.

Last but not least, the author addresses the question of algorithmic discrimination beyond protected groups. Here, two specific justifications are considered – the so-called negative spiral and arbitrariness-by-algorithm. The premise of the former is that “the scoring process generates extremely negative outcomes to some people which are disproportionate to the actual differences among the individuals” and can result in a self-fulfilling prophecy. In the latter case the problematic factor is that individuals are judged “based on what inferences and correlations suggest they might do, rather than for things they have actually done”. However, after a brief discussion, Zarsky does not find these two concerns as serious as they initially seem. He notes that, as the negative spiral problem only affects specific individuals and does so almost randomly, it does not result in a group stigma and therefore the psychological damage is not as strong. Additionally, and more importantly (he argues), the individuals harmed will generally have someone to turn to for help in stopping the downward spiral as it will not affect everyone in their social group in an equal degree. Notably, the fact that consumers are being scored in general, and offered differentiated conditions of access to goods and services as a result, does not seem to be perceived as a discrimination-based concern of its own. As regards the second issue, Zarsky observes that – assuming that the scoring process is premised upon non-spurious correlations and the errors are random and reasonable – the fact that the resulting errors are based on correlations does not make it more problematic compared to the alternatives. It is worth noting that, in their response to Zarsky, Pasquale and Citron warn against an excessive reliance on such assumptions (Pasquale and Citron, 2014).

- **Problem addressed:** Discrimination-based concerns related to scoring mechanisms
- **Key finding/argument:** Most serious concerns relate to explicit and implicit discrimination of the protected groups. With respect to disparate impact a more nuanced approach is recommended.

In some areas, such as credit and employment, the risk of segregation can be sufficiently serious to call for regulatory responses. In others, such as marketing and (tentatively) insurance, the stigma-related arguments can be more salient, yet the overall assessment will depend on a variety of factors (visibility of discriminatory processes, historical context of stigmatization). As regards discrimination beyond protected groups, the “negative spiral” and “arbitrariness-by-algorithm” arguments are considered, but eventually not perceived as exceedingly problematic from the non-discrimination perspective. The author, however, does not deny the seriousness of related autonomy-based concerns in his other works (Zarsky, 2016).

- **Solution proposed:** Throughout the paper Zarsky puts forward several policy recommendations, which generally revolve around (internal and external) auditing, licensing and education. The benefits of automated scoring (e.g. reducing human bias and, paradoxically, the stigma-reducing effects of opacity) are also pointed out. The author is more hesitant about transparency and seems to favor its milder forms such as interactive modeling. Despite the fact that, from the policy perspective, the author does not find the rationale against algorithmic discrimination beyond protected groups to be particularly strong, he does not deny the possibility of individual harm being caused in this context and considers how the respective risks could be mitigated. He advises against a requirement to carry out ex ante causation studies to establish relevancy of all factors considered by the algorithm, but rather sees a potential in ensuring that independent scoring mechanisms are applied in different contexts and that competition between different scoring systems exists. Finally, in his view, the government could be required to ensure that scoring mechanisms do not prevent individuals from fulfilling their essential social needs, such as healthcare. A similar argument is made in a later paper by the same author with respect to credit, where the idea of secure government-based loans and subsidies is advanced (Zarsky, 2016).
- Note that the paper proceeds from the assumption that the scores are premised upon individuals’ previous behavior, rely on non-spurious correlations and are followed through diligently. The author’s reliance on these assumptions may be subject to criticism (Pasquale and Citron, 2014); it shows, nevertheless, that the literature considers the elements related to the design and execution of algorithms important. To ensure such baseline conditions (e.g. procedural regularity) technological solutions, such as those advocated by Kroll et al. could be considered (Kroll et al., 2016; discussed above).
- **Related literature:** Citron and Pasquale, 2014; Pasquale and Citron, 2014.

Online sources and further readings:

- Omidyar Network, *Public Scrutiny of Automated Decisions: Early Lessons and Emerging Methods*, 27th February 2018, <https://www.omidyar.com/insights/public-scrutiny-automated-decisions-early-lessons-and-emerging-methods>
- Working Party Article 29, *Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679 (wp251rev.01)*, 13th February 2018, http://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=612053
- MIT Moral Machine Game <http://moralmachine.mit.edu/>
- Tim Sandle, *Bankers says AI will be your newest coworker*, “Digital Journal”, April 20th 2018, <http://www.digitaljournal.com/business/bankers-says-ai-will-be-your-newest-coworker/article/520324>
- O’Neill, C. (2016). *Weapons of math destruction. How Big Data Increases Inequality and Threatens Democracy*. Broadway Books.

3.3. Advertising and other commercial practices

If knowledge generation is the basis for all other uses, and automated decision making is in one way or another the nature of each use of AI by business, then **targeted and personalized advertising** is the party that seems nowadays best prepared, and where everyone wants to be. Having the ability to reach exactly the consumers who might purchase the product, in a way that they will find most convincing, is a dream come true for businesses. Big data-fueled algorithmic applications allow traders to develop better and more effective ads, by figuring out what consumers best respond to. Some consumers might be happy with this as well – but the risks of manipulation, aggressive advertising and discrimination are real.

Indeed, according to the performed market analysis and the literature overview, software agents are increasingly being used by businesses to exert technological influence upon consumers with the aim of steering them towards specific products or services. Such influence can be exerted whenever a consumer uses a connected device – be it a laptop, a smartphone, a fitness band or a smart speaker – and is closely linked to the concept of personalization.

Personalization in this context can be understood as a tailored display of information (Mik, 2016), made possible by extensive data gathering and knowledge building. Initially, the term “personalization” has been used as a slight overstatement, perhaps a marketing slogan in itself. Indeed, a marketing experience can hardly be described as “personal” if it consist in the matching of (pre-formulated) commercial messages to consumers identified as more likely to respond to these messages based on several segments (profiles) to which they were assigned. The more recent developments in data collection and analysis have, nevertheless, made it possible to optimize ad exposure to a much higher degree. The development is presented to consumers as a source of an added value – after all, the advances in microtargeting increase the relevance of commercial messages addressed to them. Unsurprisingly, however, concerns about consumers’ wellbeing are also voiced (Mik, 2016; Calo, 2014).

Two factors are particularly important for this context: 1) the intransparent, diverse and potentially highly persuasive nature of the different “nudges” (Thaler and Sunstein, 2009), deployed by commercial actors in the digital sphere (e.g. ranging from the order of search results to the specific commercial messages addressing a consumer) and 2) the novel kind of information asymmetries arising between suppliers and consumers (see also section 1.3.1 above). Of course, the very fact that consumer decisions are being influenced by advertisers, who resort to a wide variety of techniques, is in itself nothing new. Both in the online and in the offline world “advertising”, in the strict sense of the word, is by far not the only tool of exerting influence on consumer decisions (take the example of the layout of a supermarket and the order of the search results). It has been argued, nevertheless, that the methods of “choice

architecture” deployed by businesses online can be radically more effective compared to their offline counterparts. As noted by Mik, the scale, variety and effectiveness of the technology-enabled influence are not comparable to any marketing strategies known from the offline world. In her view, it is especially the “combined, mutually-enforcing effect of multiple technologies” that raises questions about the “legally permissible levels of transactional exploitation” (Mik, 2016). Along similar lines, Calo warns against the threats of “digital market manipulation” (Calo, 2014).

Indeed, the analytical and predictive techniques that can nowadays be applied to create granular pictures of particular consumers and interact with them provide the traders with the different, more powerful tools of exerting influence (also referred to in the literature as “hypernudges”; Yeung, 2017). Most notably, the insights obtained about consumers are no longer limited to their established or inferred traits and preferences, but also include their (mis)perceptions and vulnerabilities, such as psychological and physiological states (Bar-Gill, 2018; Helberger, 2016; Mik, 2016). With further advances in big data analytics and the growing uptake of the Internet of Things, the accuracy of insights obtained about particular consumers is bound to increase even more. This suggests that the transformation brought about by the digital technology is not merely a one of degree. Although the knowledge that digital businesses possess about consumers is still largely based on correlations and segmentations (the latter becoming more accurate and granular), arguably, a substantive change is also taking place. A growing number of reports focus on businesses’ ability to obtain *genuinely individual* insights about particular consumers (a phenomenon which is sometimes described as individualization⁴³ or extreme personalization⁴⁴). This allows not only for a more effective decision-making regarding *exposure* to commercial communications, but also the *creation* of advertisements and products tailored to specific consumers, with all the related concerns discussed before.

Use type	Advertising and Other Commercial Practices
In what way is business employing AI for this use?	<ol style="list-style-type: none"> 1. Ads creation: testing how well particular groups of consumers respond to different types of commercials; improving the effectiveness of different types of communication 2. Ads delivery: Displaying ads to those who have the highest probability of clicking/purchasing the product/ service 3. Targeted offers: fine-tuning the characteristics/price of a product/service to the needs of a given category of consumers, or even a particular consumer
What is the gain for business?	<p>Ad 1 More effective content of ads</p> <p>Ad 2 More effective distribution of ads</p> <p>Ad 3 Higher chance of concluding a transaction</p>
What is the gain for consumers?	<p>Potentially more relevant content displayed</p> <p>Potentially lower transaction/search costs</p>

⁴³ <https://www.targetmarketingmag.com/article/individualization-new-personalization-content-marketing-resolution-2018/all/>

⁴⁴ <https://www.forbes.com/sites/briansolis/2017/11/30/extreme-personalization-is-the-new-personalization-how-to-use-ai-to-personalize-consumer-engagement/#51ee3dc9829a>

What is the loss for consumers?	<p>Risk of aggressive advertising (exploiting consumers' vulnerabilities)</p> <p>Risk of discrimination in advertising</p> <p>Risk of seeing only the content falling into historically revealed preferences (risk of missing out)</p>
Third party effects/ externalities	<p>Growing dependence on targeted ad delivery providers</p>
Other	<p>The distinction between an advertisement and an offer, an add and a contract, is getting more and more blurred. A more in-depth typology of commercial practices, as they currently are, might be needed</p>

Discussion of selected papers

Draper, N. A., & Turow J. (2017). Audience Constructions, Reputations, and Emerging Media Technologies: New Issues of Legal and Social Policy. In: R. Brownsword, E. Scotford & K. Yeung (Eds.), The Oxford Handbook of Law, Regulation and Technology. Oxford University Press.

Draper and Turow argue that the phenomenon of audience construction in advertising – the will to reach a particular category of consumers by advertisers – is nothing new and can be traced all the way back to different media, including television, radio and press. Advertisers always tried to differentiate who sees what commercial communication by, among others, choosing to place ads for women/men/children in magazines culturally associated with this audience; and used to deploy the categories currently informing different “profiles” long before the age of AI. What has changed now is the scale on which this can be done, both regarding the amount of audiences (much finer profiling), and accuracy of actually reaching them. For the purposes of this project, what matters is that: profiling is not a new phenomenon, on the contrary – it is a natural consequence of how marketing has always functioned. Maybe involving marketing experts in research projects regarding consumer law and ads is a good idea?

Helberger, N. (2016). Profiling and Targeting Consumers in the Internet of Things – A New Challenge for Consumer Law. In: R. Schulze & D. Staudenmayer (Eds.), Digital Revolution: Challenges for Contract Law in Practice. Nomos.

The author discusses the challenges associated with profiling and targeting against the background of European consumer law. Most notably, Helberger attempts to link the observed increase in the role played by collection and use of data with consumer protection concepts and provisions such as those related to consumer information and fairness in commercial transactions. According to the author, there is already an obligation on the part of the traders to inform consumers about the use of tracking techniques and personalized communications under the existing consumer law (Directive 2011/83/EU on consumer rights). As regards the problems of automated decision-making, the author considers that this could be regarded as an indicator of contractual unfairness – where the sharing of data “has the potential to lead to consumer detriment, algorithmic discrimination or other forms of unfavourable decision-making”. The same could also be true, in her view, for certain forms of “nudging” or personalized advertising, especially when consumer choice is limited due to trader’s monopoly or when the consumer is particularly weak. The source of one of such weaknesses Helberger sees in the consumer’s susceptibility to the “digital market manipulation” (Calo, 2014). The subsequent part of the

chapter discusses the idea of digital market manipulation in more detail and points to the potential of Unfair Commercial Practices Directive in countering its negative effects.

- **Problem addressed:** The challenges of profiling and targeting which affect not only privacy and data protection, but also the protection of contractual fairness, adequate information and autonomous and free consumer choices.
- **Key finding/argument:** Consumer law, and in particular Unfair Commercial Practices Directive, can be treated as a point of departure for the conceptualization of fairer marketing practices in the digital age (here, particularly, in the Internet of Things). In playing this role, the unfair commercial practices law can complement the data protection framework.
- **Solution proposed:** Further elaboration of the UCPD notions.

Online sources and further readings:

- Julia Carrie Wong, 'It might work too well': the dark art of political advertising online, "The Guardian", March 19th 2018,
- <https://www.theguardian.com/technology/2018/mar/19/facebook-political-ads-social-media-history-online-democracy>
- Latanya Sweeney, Discrimination in Online Ad Delivery, 2013, <https://arxiv.org/abs/1301.6822>
- Daniel Fagella, Artificial Intelligence in Marketing and Advertising – 5 Examples of Real Traction, "techemergence", January 17th 2018, <https://www.techemergence.com/artificial-intelligence-in-marketing-and-advertising-5-examples-of-real-traction/>
- Alex Hern, Can we really trust Google as judge, jury and executioner of online ads?, "the Guardian, February 2015, <https://www.theguardian.com/technology/2018/feb/15/google-chrome-adblocking-online-ads>

3.4. Personal Assistants ("Algorithmic Consumers", "Cyber-Butlers")

On May 8th 2018 Google demonstrated the new capabilities of the Google Assistant, which (who?) can now make phone calls on behalf of a user, in order to book a reservation at any small business, like a restaurant or a hairdresser's. What thrilled the audience and perplexed many commentators was not only the Assistant's impressive ability to achieve its task in the light of unpredictable responses of a human interlocutor, but most of all *how closely it imitated the human way of speaking*. With reactions like "um" or "mhm", voice modulation and pauses, it was close-to-impossible to tell that the phone call was being held by a machine. Quickly labelled "deceitful AI", the technology's new capabilities gave rise to questions about whether a robot should always disclose its robotic identity, and whether we are getting close to passing the Turing test.

The debate is by no means new, though. Personal digital assistants, coming under the labels of "cyber-butlers" (Brownsword 2017), "algorithmic consumers" and others, have been attracting the attention of scholars and media for quite a while now. Apple's Siri, Google's Assistant and Amazon's Echo feature in countless articles and videos online, ranging from voices full of praise, to those filled with concerns. The most advanced, commercially available personal assistants are currently based on automatic speech recognition, natural language understanding and text-to-speech technologies. Agents of this kind can be used to "voice-control one's world", for example to make and handle phone calls (e.g. to make an appointment), place orders in online stores, manage one's agenda, play music, check traffic or weather conditions, read the news, control one's smart home, etc. They are available on a variety of mobile and smart home devices, such as smartphones and smart speakers. Other, usually text-based agents (chatbots) can also be accessed via messaging applications such as Facebook Messenger or WeChat.

On the one hand, commentators underline that these devices – existing also as a part of the Internet of Things, connected to one another and initiating communication – make consumers' lives significantly

easier. They allow them to make purchases, retrieve information, make their homes smarter etc. without much effort. Some of the functions performed by digital assistants are, in themselves, a response to challenges created by technology. One of such examples is the use of personal assistants to make recommendations, which allows consumers to navigate through the exponentially growing data volumes (software agents are able to dramatically reduce search costs). Attention can further be drawn to the idea of “algorithmic buying groups” (Gal and Elkin-Koren, 2017), which can, at least theoretically, be formed on the demand side not only to negotiate better conditions, but also to reduce privacy- and discrimination-based concerns associated with knowledge building and automated decision-making, as discussed in previous sections of this report. A similar idea is also reflected in less advanced solutions such as credit report products, which have long been available on the market. Products of this kind offer consumers access to their own credit history and score, and provide advice that can be used to improve one’s score and minimize interest rates (Langely, 2014). It is not entirely unthinkable that, in a not-too-distant future, advice of this kind could be provided by artificial agents (also referred to as “robo-advisors”; see Lightbourne, 2017).

At the same time, virtual assistants also pose a number of economic and non-economic risks for consumers. Some of them are rather anecdotal (and, in the literature, usually discussed by contract lawyers), such as the risk of transactions performed as a result of machine errors or compliance with a child’s request. Other, like those related to unauthorized access by third parties and systems’ critical failures (security-related concerns), constitute crosscutting risks associated not only with algorithmic consumers, but with software agents in general. Naturally, in both contexts, the degree of harm suffered by consumers will depend according to the context (the harm caused by a breach to consumer’s smart home may be more serious compared to a situation when third party manipulation affects the content of displayed ads).

More importantly, personal assistants allow their suppliers to gain access to and understanding of a vast amount of consumer data. This is due to the personal assistants’ ability to “hear everything”, since their microphones and other input devices are always “on”. The fact that they often communicate with smart home (data from physical sensors) and smartphone (data from virtual sensors) exacerbates the problem. This results in privacy-related concerns of an unprecedented degree as well as in the equally significant risks of consumer cognitive limitations and vulnerabilities being exploited for economic gain. The latter is particularly true when artificial agents represent a variety of different interests at the same time, of which consumers are not aware. By way of illustration, if artificial agents are deployed by businesses within their internal structures, consumers should generally be aware of the need to evaluate their operations critically. This can be the case, *inter alia*, for chatbots used to facilitate customer service or the more advanced internal robo-advisors, particularly if active in the sectors where they cannot be expected to act as fiduciaries. In these situations the impact of digital assistants on consumers’ decision-making capacities might not be a prime concern (other issues may, in turn, become more salient, like the consumer’s ability to effectively communicate with the company and bring his or her claims). The same, however, is not necessarily the case for intelligent assistants provided in the remaining contexts mentioned above, including in the most relevant case of personal assistants provided by companies like Google or Amazon. These products are, nevertheless, far from the idea of a “virtuous assistant”, in which consumers’ interests are prioritized by default. Instead, they rather involve an ambiguous and opaque network of interests, which puts further strain on the already fragile boundary between advertisement and information. Following Baldwin, the initial question asked here is whether the context in which consumers’ decision is subject to an external influence is manageable (Baldwin, 2014). Where such an influence is not perceived, consumers’ ability to neutralize its effects becomes limited, which, in turn, poses a risk to his or her freedom of choice.

A related challenge concerns the potential decrease of the range of options effectively available to consumers. This results not from the economic calculations mentioned above, but from purely technical and unintended factors, such as an inability of an algorithm to detect certain nuances in the consumer profile or the incorrect assumptions having been made and reproduced. The risks voiced in this regard

concern: a) monopolization/de-neutralization of distribution channels in consumer markets (designing purchasing systems in a way that makes them prefer the products offered by the same corporate group, or “trusted”, i.e. paying, partners); b) lower quality of products/services, or c) a mismatch between consumers’ actual and revealed/inferred preferences.

Yet another point related to consumers’ autonomy refers to the very basic idea of outsourcing decision-making powers to artificial agents. Admittedly, one could argue that this dimension of digital assistants is not problematic from the perspective of autonomy and dignity, as the original autonomous decision to rely upon such an agent can be attributed to the consumer. This is, nevertheless, only half of the story. The argument is illustratively described by Brownsword:

At first blush, the use of such PDAs [personal digital assistants] might seem to be unproblematic. If Diana chooses to rely on her PDA, how does this challenge human dignity? However, if it turns out that Diana’s freedom to use her PDA is more apparent than real, her ‘consent’ might not be adequate to authorise the kind of exploitation of her personal data that might trouble advocates of HD2a [the rights-based conception of human dignity]. (Brownsword, 2017)

Arguably, this part of Brownsword’s analysis can equally be applied to the discussion of “private autonomy” in the civil law understanding. Here, a question may indeed be asked about which conditions must be fulfilled for a consumer’s representation to become legally binding. Moreover, when the subsequent decisions of the PDA exploit the consumer’s cognitive limitations and behavioral biases, a broader normative assessment might be required of the “legally permissible levels of transactional exploitation” of a systemic nature (see also Mik, 2016).

The dignity-based concerns discussed by Brownsword do not stop here, however. As further argued by the author:

More importantly, we need to take a harder look at the way in which Diana relies on her PDA. In particular, what if Diana comes to rely on her PDA to analyse her moral dilemmas and determine how she should act? For those who belong to an aspirant moral community, it is axiomatic that each agent should develop a sense of what it is to do the right thing for the right reason, and try always to do just that. Human dignity involves more than merely acting in line with the right thing; the paradigmatic expression of the dignity of humans is in doing what an agent judges to be the right thing even where there is an opportunity to do the wrong thing. (Brownsword, 2017)

The autonomy and dignity-based concerns thus do not appear to be less pronounced when personal assistants are employed by consumers to perform different tasks, compared to a situation in which a consumer is clearly an object of an automated decision. As a matter of fact, the contrary can be true. This is particularly the case when the market for PDAs becomes dominated by a limited number of mega platforms, which further a number of interests that cannot be aligned.

It is therefore not surprising that the existing literature focuses on the role of such mega platforms in the emerging markets for digital assistants. The prospects flowing from this assessment are rather bleak. According to Stucke and Ezechia (2017), at least three key elements bode in favor of such mega-platforms taking over the personal assistant’s market and keeping its main determinants unaffected: data, network effects, and the scope of platform’s services (or else “tying free services with algorithmic consumer functions which build upon economies of scale, scope and speed” as noted by Gal and Elkin-Koren, 2016). This poses a risk to consumer welfare, for example when there is a significant mismatch between consumer purchases and consumer preferences. Along with the aforementioned considerations of imperceptible influences, the described situation leads to an intuitive call for more transparency. However, consumers’ inability to assess the extent to which the respective risks may materialize would likely not be overcome by traditional rules on mandated disclosure. When trying to understand the choices made by algorithms after being provided with relevant information, consumers would continue to face the very same cognitive problems as the ones that algorithms seek to counteract in the first place (Gal and Elkin-Koren, 2017). Similar factors are also believed to prevent consumers from switching, even if competing and more privacy-friendly products are available (as illustrated by the market of search engines; Stucke and Ezechia, 2017). On top of this, new cognitive biases are also, arguably, being

created – or amplified – by the digital transformation itself. An example of this could be the perceived tendency to take authenticity of digital data for granted, as evidenced by the alarming expansion of fake news.

Use type	Digital Personal Assistants
In what way is business employing AI for this use?	<ol style="list-style-type: none"> 1. Sales of digital assistants themselves (selling products) 2. Enabling “smart solutions” in the environment of consumers 3. Collecting information about the preferences and habits of consumers 4. Facilitating communications with consumers through chatbots 5. Means of advertising/suggestions 6. Delivery of requested and/or targeted content/goods and services
What is the gain for business?	<p>Ad 1/2 Direct profit</p> <p>Ad 3 Source of data about consumers preferences and behaviours</p> <p>Ad 4 Faster and cheaper communication</p> <p>Ad 5/6 Additional channels for (personalized) commercial practices</p>
What is the gain for consumers?	<p>Ad 2 More convenient home solutions</p> <p>Ad 4 Potentially easier and faster communication with businesses</p> <p>Ad 5/6 Easier access to content, goods and services</p>
What is the loss for consumers?	<p>Ad 1/5/6 Growing dependence on the digital personal assistants</p> <p>Ad 2 Security risks in case of hacking and/or systems’ failure</p> <p>Ad 3 Privacy concerns, both dignitarian (businesses know more and more about consumers, including sensitive data) and consequentialist (risk of discrimination, aggressive commercial practices)</p> <p>Ad 4 Potentially lower quality of information/more difficult communication</p> <p>Ad 5 Potentially discriminatory/aggressive/otherwise exploitative content</p> <p>Ad 5/6 Decrease in the autonomy and freedom of choice of consumers</p>
Third party effects/ externalities	<p>Potential job losses (salespersons, call persons)</p> <p>Risks of unfair competition for SMEs</p> <p>Higher threshold of market entry</p>
Other	<p>The general philosophical question one faces here is the choice between freedom and privacy on the one hand, and convenience and efficiency on the other</p>

Discussion of selected papers

Gal, M. S., & Elkin-Koren, N. (2017). Algorithmic Consumers. *Harvard Journal of Law & Technology*, 30(2), 309.

The article investigates the impact of algorithmic consumers, i.e. digital agents designed to facilitate or handle commercial transactions for consumers. It presents both benefits and challenges of algorithmic consumers. Overall, it appears to place more emphasis on the benefits and focuses on the role of antitrust law in that regard.

Problem addressed

At the individual level, the identified problems include: the possible reduction of consumers' autonomy and choice as well as an increased vulnerability to certain consumer harms caused by unauthorized access and manipulation. The authors also refer to the broader societal implications, such as the potential loss of social interactions and the growth of inequality between designers/owners on the one hand, and consumers/suppliers on the other. They note, at the same time, that the deployment of algorithmic techniques by consumers can also affect these factors in a positive way. Most notably, as noted by the authors, algorithmic consumers are capable of significantly enhancing consumer choice. They can also make consumers less vulnerable to some risks. For example, the fact that algorithms are immune to bias can render targeted online advertising less relevant (which, in itself, can be considered both as a merit and a flaw). Furthermore, the operation of "algorithmic buying groups" can reduce the risks related to privacy and discrimination.

Key finding/argument

At the market level, the deployment of algorithmic techniques by consumers can lead to a possible increase in consumer welfare, lowering the costs and providing better quality. However, the extent of this increase will depend on several factors, most notably the market power of the algorithm providers, relative to (a) suppliers (stronger algorithms provide stronger incentives to compete on the merits), (b) consumers (possibly less value passed on to consumers if the market power of algorithm providers is strong), (c) competing algorithm providers. The competitive pressure placed upon algorithm providers depends, in particular, on entry barriers to algorithmic markets as well as, to some extent, on the consumers' ability to compare competing algorithms. The main identified barriers to entry include: access to data (input market), access to consumers (output market) and exclusionary conduct (excluding competitors – e.g. by entering into exclusive dealings contracts with suppliers or excluding suppliers by choosing not to buy from them).

Solution proposed

The solutions proposed focus on the regulatory tools addressing the potential entry barriers into algorithmic consumer markets. The perspective is thus significantly different from that of the present project. Some links can, nevertheless, be established.

Attention should be drawn, in particular, to the solutions related to "reducing barriers to consumer access" (understood as access to consumers). These barriers may arise from many factors, including imperfect information on the part of consumers. If the consumer's information is limited, then the consumer is less likely to switch providers. Consequently, it seems that the tools aimed to enhance consumer decision-making capacities can also be conceptualized as aimed at the reduction of the competitive barrier to consumer access (note that this point is not made by the authors).

In this regard, the authors note that the tools mitigating this barrier can either be provided by the market (e.g. comparison firms) or have a regulatory nature. As regards the latter, instruments mentioned by Gal

and Elkin-Koren include the laws prohibiting misleading information and the laws requiring transparency about some product features. However, these elements are not discussed in more detail (“In this regard algorithms are no different from other products [except for the black box effect]” – there is no discussion of the black box, however). This is probably because the authors prefer to focus on what they perceive to be a more significant barrier, namely the access to consumers via intermediaries (mega platforms like Google, Apple, Facebook, which are also most likely to develop algorithmic consumers). The scholars subsequently attempt to tackle this concern from the antitrust angle, but eventually acknowledge the limits of the (US) antitrust law as it currently stands. Therefore, the challenge regarding the “erection of entry barriers that arise from the tying of free services with algorithmic consumer functions which build upon economies of scale, scope and speed” remains.

With respect to access to data, the role of antitrust law is again discussed and its limitations are similarly acknowledged. According to the authors, these barriers can, to some extent, be reduced by the market – for example when data sources alternative to mega platforms are offered by the IoT. Other complementary tools, briefly addressed in the paper, could encompass the rules on data portability and interoperability.

Finally, an aspect, which the paper addresses in most detail, relates to “exclusionary conduct”. The focus once again remains on the role of antitrust law, about which, by contrast, the authors appear to be more optimistic. More specific solutions within this domain are also offered (e.g. how to assess whether parallel conduct amounts to an agreement, how to assess which agreements are anticompetitive, what is an anticompetitive intent).

Online sources and further readings:

- Chris Welsh, *Google just gave a stunning demo of Assistant making an actual phone call*, “The Verge”, 8th May 2018, <https://www.theverge.com/2018/5/8/17332070/google-assistant-makes-phone-call-demo-duplex-io-2018>
- Alex Cranz, *Google Assistant Impersonating a Human Is Scary as Hell to Me*, “Gizmodo”, 9th May 2018, <http://www.gizmodo.co.uk/2018/05/uhh-google-assistant-impersonating-a-human-is-scary-as-hell-to-me/>
- Alex Harn, *Google's 'deceitful' AI assistant to identify itself as a robot during calls*, “The Guardian”, 11th May 2018, <https://www.theguardian.com/technology/2018/may/11/google-duplex-ai-identify-itself-as-robot-during-calls>
- Lance Ulanoff, *Did Google Duplex just pass the Turing Test?*, Medium.com, 8th May 2018, <https://medium.com/@LanceUlanoff/did-google-duplex-just-pass-the-turing-test-ffcfe6868b02>
- Adam Dechis, *The Seven Best Things You Can Do With an Amazon Echo*, “Life Hacker”, 30th March 2016 <https://lifehacker.com/the-seven-best-things-you-can-do-with-an-amazon-echo-1766989219>
- Rory Carroll, *Goodbye privacy, hello 'Alexa': Amazon Echo, the home robot who hears it all*, “The Guardian”, 21st November 2015 <https://www.theguardian.com/technology/2015/nov/21/amazon-echo-alexa-home-robot-privacy-cloud>
- Lily Hay Newman, *Turning an Echo into a spy device only took some clever coding*, “Wired”, <https://www.wired.com/story/amazon-echo-alexa-skill-spying/>

4. Time for Phase Two: Further questions in place of conclusions

The very aim and ambition of this project have been to pave the way for the next stage of interdisciplinary research. We departed from the premise that there is not **one** AI problem, and there will be not one solution to challenges posed by AI. On the contrary, the response will be different across the sectors, both regarding the substance and form of the potential regulatory response. In the previous sections of this report we have provided an overview of the ways in which businesses employ artificial intelligence in relations with consumers, as well as analyzed what the risks and the gains are for whom, stemming from that use. We have gathered empirical material and reviewed the relevant scholarly literature in order to shed more light on the respective risks and gains for consumers and evaluate them through the lens of consumer law principles. Now, it is time to ponder what to do about all this.

The research undertaken throughout this project was supposed to serve as a starting point for the further analysis in the concrete, interdisciplinary research environments. We hope that it can serve as a valuable source of inspiration, knowledge and orientation for scholars interested in pursuing this path further. In this section we present 25 research questions, in a seed form. We believe that every single one of them is important and must be addressed not only by lawyers, but by interdisciplinary teams comprising engineers, data scientists, normative scientists, and governance experts. Only in a mutual dialogue will we be able to move forward. We group the questions into four categories: 1) Normative (what should be the goals we strive for?); 2) Technological (how is it technologically possible to achieve these goals?); 3) Governance (how should the relevant legal frameworks and the broader incentives systems be designed) and 4) Theoretical. As we claim, *all of them*, all of the types, should be pondered jointly – feasibility (technical and political) influences the choice of goals, the formulation of goals influences the range of possible means. However, we believe there is a value in presenting them separately.

4.1. Normative questions (goals)

Question 1: Should there be limits to price discrimination in consumer retail markets? If so, what limits: type of good/service, type of criteria?

With targeted communication based on profiling and individual data, businesses are able to offer the **same goods** and services for **different prices** to different customers on a **mass scale**. As of now, “price discrimination” is not forbidden by EU consumer law⁴⁵. Even more, one could argue that the ability to offer different prices is one of the foundations of the market economy, based on bargaining, individual profit seeking and negotiations. On the other hand, the standard practice in consumer retail markets, with which in mind consumer law had originally been developed, has been such that everyone pays the same price for the same publically offered goods and services. If a millionaire and a beggar walk into McDonald’s, they will both pay \$1 for their cheeseburger. Everyone pays the same price in a supermarket, displayed on the product/shelf. Did this practice already give rise to a legal standard? The reason for that practice is not businesses’ egalitarian mindset, it is (was) just the cheapest to offer the same retail price to everyone who walks in, without engaging in negotiations with everyone. Now, with targeted offers on publically available websites (and, at some point, possibly also retail stores), based on profiles and individual data, quickly computed by robust machines, every consumer can be displayed a price corresponding to what the system assumes is the highest number that they are willing to pay. A student from a poor country can see a lower price for a flight than a professor from a rich country.

Maybe this is even desirable? Could price discrimination be to the benefit of consumers, when the poor pay less and the rich pay more? Or maybe there is a risk that price discrimination will actually make the

⁴⁵ See e.g. European Commission, Guidance on the Implementation/Application of Directive 2005/29/EC on Unfair Commercial Practices, SWD(2016) 163 final, p. 134.

poor pay even more? What if the businesses rely not only on consumers' willingness to pay, but also on their misperceptions (vulnerabilities, cognitive biases) or personal characteristics (e.g. race, age, gender) in setting respective prices? Is it problematic that prices are set and disclosed bilaterally and cannot easily be compared ("do we still have a market")? Should price discrimination be encouraged/discouraged or controlled in specific markets or sectors, e.g. when the goods are (not) easily substitutable or with universal services (water, electricity, different media)? Should scoring and algorithmic decision-making in cases of credit etc., when costs borne by traders are heterogeneous, be treated differently from products with homogenous costs of supply? Should we rethink the notion of substitutability? Could the principle of weaker party protection justify a more hands-on policy on personalized pricing – especially if we conclude that the principle assumes particular importance in the digital age, due to the new and unprecedented asymmetries between traders and consumers? What about the principle of regulated autonomy and the consumer's freedom to choose as one of its main components? Is there a risk that consumers will get "stuck" in their profiles? Could personalized quality be the next big question?

Recommended readings:

Bar-Gill, O. (2018). Algorithmic Price Discrimination: When Demand is a Function of Both Preferences and (Mis)perceptions. *University of Chicago Law Review*, Forthcoming

Question 2: Should discrimination in consumer markets be forbidden? If so, under what conditions and what does it mean for algorithms?

In the EU there is a body of laws that prohibit discrimination in access to good and services. However, as for now, one cannot speak of a comprehensive set of rules and even when such rules exist the logic behind them is not always consistent. Some of them appear to be more concerned with the internal market building, while others seem to focus on the protection of persons with specific characteristics such as racial or ethnic origin, gender or sexual orientation from a negative socio-economic impact (segregation, stigmatization). A question then becomes warranted: can we agree whether discrimination in access to goods and services should be prohibited and, if so, based on what kind of rationale? Practice shows that a normative assessment can be difficult to make even in the most clear-cut cases of different treatment on the basis of concrete personal characteristics, e.g. when a gay couple is denied a wedding cake. However, the question only becomes more complicated as one goes deeper. For example, scholars have long pondered upon the complex questions of affirmative action and intersectionality. It might be desirable to take certain protected characteristics into account to right the wrongs of the past or prevent the negative consequences from occurring. And what if a combination of different characteristics creates a cumulative disadvantage? Should the law provide additional responses and, if so, what should be the respective parameters and goals? What about economic factors – are they entirely irrelevant for this debate? Now add algorithms to the equation, as already hinted at in Question 1. Is it problematic if the factors that software agents take into account in producing their outcomes (e.g. deciding on consumer's eligibility for the product and/or presenting an offer) include personal characteristics like ethnic origin or gender? If so, is it problematic *per se* or only when the negative outcomes materialize?

Recommended readings:

Zarsky, T. (2014). Understanding Discrimination in the Scored Society. *Washington Law Review*, 89, 1375.

Question 3: Should responses to the issue of discrimination in consumer markets differ according to the context? For example, does discrimination resulting from the use of algorithmic techniques in specific markets – e.g. to assess credit eligibility or calculate insurance rates – call for targeted responses?

Certain consumer goods and services have traditionally been subject to more regulation than others due to the importance and/or risks they carry for consumers. Financial markets are a prime example. Are these also the markets where questions of differentiated treatment should be scrutinized in most detail? Indeed, the issue of algorithms used in particular by credit institutions to score consumers and assess their eligibility for particular services have drawn a great deal of scholarly attention. Traders' incentives to profile consumers in such markets are understandable from an economic point of view, considering the heterogeneity of costs associated with specific transactions. However, are there any red lines? What if credit conditions offered to women are consistently less favorable than those offered to men? Does it matter where the source of this outcome lies or what the trader has done/could have done to prevent it? For example, does it matter whether a discriminatory outcome follows from an application of a machine-learning algorithm? Often non-discrimination laws refer to "objective criteria" / "business reasons" as a possible justification for (indirect) discrimination. What if algorithms become "too good" in understanding the reality and their discriminatory predictions are, in fact, correct? Should businesses be obliged to adjust their conduct nevertheless or should action rather be expected from public regulators and the civil society?

Recommended readings:

Barocas, S., & Selbst, A. D. (2016). Big Data's Disparate Impact. *California Law Review*. 104, 671.

Citron, D. K., & Pasquale, F. (2014). The scored society: due process for automated predictions. *Washington Law Review*, 89, 1.

Helveston, M. N. (2016). Consumer Protection in the Age of Big Data. *Washington University Law Review*, 93(4), 859.

Question 4: Should discrimination in online ad delivery be forbidden? If not, should there be limits?

It is no wonder that there is considerable public attention for questions of discrimination in assessing the eligibility of consumers for a product or service that they seek to access in order to fulfill their core societal needs. It seems intuitively wrong to make it impossible for them to access such goods or services, or to do so only on prohibitive conditions, purely because of their race or similar characteristics – irrespective of where the source of the problem lies. But what if such differentiated treatment applies to marketing messages? With targeted advertising based on profiling and individual data, businesses are able to display advertisements, which are highly "personalized", including discriminatory. Following a long-established aim to show commercials to those who might want the product most (toys commercials on TV channels with cartoons, private banking commercials on business news channels, make-up advertisements in magazines aimed at women, etc.), businesses are now able to craft the choice for every single consumer out there. There are benefits, of course: if one *must* see ads (that is why so many online services are "free", at least in monetary terms), one might just as well see *interesting* ads. However, there are also challenges. Businesses maximize the amount of clicks, and the personalization is often based on historical data, which might be biased. As a result, as studies have shown, commercials for high-paid jobs tend to be displayed to white males, while criminal records bureaus tend to pop-up when googling an African-sounding name.

As of now, discrimination in ad delivery – unlike discrimination in access to goods and services (even if only partially) – is not forbidden. However, this might be due to the fact that such discrimination was simply impossible before the advent of the algorithmic era. There is a difference between choosing an

audience (where someone might still “break in” – no one stops a poor student from reading a business magazine), and choosing a specific client (who has no ability to “change” his profile). On the other hand, one can imagine instances of differentiation of ads based on categories like race, religion or gender, which seem normatively acceptable. Consider showing books on catholic theology to Catholics, make-up foundations of a particular color to people of a particular skin tone, or clothes corresponding to cultural norms associating them with a concrete gender to people of that gender. This seems to be less problematic than differentiating based on these criteria when advertisements of jobs, or housing, or all-inclusive package travels are concerned. Should discrimination in ad delivery, based on certain criteria, be forbidden? As a rule with exceptions or as an exception to the rule? Should it be process-based or outcome-based? If the focus remains on the process, how much efficiency are we willing to sacrifice, given that a lot of data might indirectly indicate protected groups? If the focus remains on the outcome, how much is too much? If 51% of a job’s ad audience is white male, is that discrimination already? 60%? 90%?

Recommended readings:

Chander, A. (2017). The Racist Algorithm?. Michigan Law Review, 115, 1023.

Question 5: Should consumers be protected from targeted advertisements and offers that seek to exploit their vulnerabilities?

There is a thin line between showing me what **I need**, showing me what **I think I need**, and showing me what **I fear I need**. With highly personalized advertising, businesses might coerce consumers into commercial decisions they would not have taken without the circumstance of “vulnerability”. *Everyone* is vulnerable every now and then. A person suffering from insomnia might buy sleeping pills when awake at 3am, someone constantly looking for a job might purchase a sketchy guide on how to prepare for job interviews, etc. As of now, Unfair Commercial Practices Directive 2005/29/EC, defines commercial practice as aggressive when “*in its factual context, taking account of all its features and circumstances, by harassment, coercion, including the use of physical force, or **undue influence**, it significantly impairs or is likely to significantly impair the average consumer’s **freedom of choice or conduct** with regard to the product and thereby causes him or is likely to cause him to take a transactional decision that he would not have taken otherwise*”. With the access of big data, merged with the insights from behavioral and cognitive studies, businesses currently understand consumers’ decision-making behavior better than consumer themselves. What is more, it is highly possible that an artificial agent *itself* learns how to exploit vulnerabilities, from the thousands of available data points and without anyone’s supervision or direct command, even if its task has been neutrally defined as “display the ad to X persons with the highest possibility of clicking, at the moment when the chance is the biggest”. One could argue that if consumers fall for these types of tricks, it is their fault, and the law should not interfere. One could just as well argue, however, that the role of consumer law is to intervene to protect weaker party, particularly at the times when he or she is vulnerable.

Recommended readings:

Calo, R. (2014). Digital Market Manipulation. George Washington Law Review, 82, 995.

Sax, M., Helberger, N., & Bol, N. (2018). Journal of Consumer Policy, 41, 103

Question 6: Should the use of certain (more invasive) profiling techniques be forbidden? If not, should there be limits?

The risks mentioned above can generally be linked to the traders’ use of algorithmic techniques for generating knowledge about consumers and acting upon that knowledge. The same goes for the further set of risks, which relate to the *influence* that traders can exert on consumers’ decisions thanks to the

use of such techniques. An underlying concern here is that the action that businesses can take after the relevant knowledge has been generated can become “too effective”. In other words, the use of algorithmic techniques by traders allows them to maximize their benefits without much regard to the various consumer interests (charge maximum prices consumers are willing to pay, present targeted messages when consumers are most vulnerable, engage in discriminatory treatment). As discussed before, one may consider such outcomes to be unwarranted and seek for regulatory responses.

One may, however, also look at the ways in which the knowledge has been *generated* in the first place and find it desirable to already put limits there. A similar logic underlies the EU data protection law, which focuses on the processing of personal data. Following the spirit of this framework one may wish to place the limits on the categories of data being processed (see Article 9 of the GDPR on the processing of sensitive data) or on the means used to collect a given piece of (personal) data (see the provisions on “cookies” in the ePrivacy Directive 2002/58/EC, currently under revision). In both cases emphasis is placed on the data subject’s explicit consent. In practice, however, consent often remains theoretical – either due to consumers’ cognitive limitations and behavioral biases or due to its effective tying to the use of the product or service. If we agree that there should be limits to the ways in which knowledge about consumers can be generated, does consent in general as well as its current legal framework present a viable solution? Or should certain ways of acquiring knowledge – for example advanced facial recognition technologies – remain off limits to the business?

Recommended readings:

Pearce, H. (2015). Online data transactions, consent, and big data: technological solutions to technological problems?. *Computer and Telecommunications Law Review*, 21(6), 149

Yeung, K. (2017). ‘Hypnudge’: Big Data as a mode of regulation by design. *Information, Communication & Society*, 20(1), 118.

Question 7: Should consumers have a right to “turn” personalization “off” or be able to influence what content they are being displayed?

Building upon the previous question, one could ask to what extent consumers could be given a more general right not to be profiled and personalized, with benefits and losses of opting out. Seeing commercials that one could find interesting may be more pleasant than seeing commercials of completely random products; getting a lower price is always nicer than getting a higher one. Yet targeted advertising could also expose consumers to a risk of discrimination/aggressive communications and does not exclude the possibility of the consumer ultimately paying more. What if a consumer wants to remain an “average consumer”? Should one have such a right? Should it be general or specific? If specific, should it apply to certain types of markets/data? How to make such a right operable, considering past experience with consent? Consider the idea behind the recent regulation 2018/302 on geo-blocking, according to which consumers should generally be able to look at the conditions at which a product or service is offered to consumers from other Member States⁴⁶ and (in a more limited number of cases) make purchases on these terms without being hindered from doing so. Could this idea be generalized (e.g. should it be possible to buy online anonymously)? What would a non-profiled commercial communication even mean? It is probably fine to show ads in a given language in a jurisdiction where this language is most spoken; e.g. to display offers of academic books on an academic website. Should the limit of personalization be “the foreseen audience of the place”?

Alternatively, should consumers be given a possibility to influence the type of content that is being displayed to them: either by resorting to technological tools, or by being provided with information on how to adjust their conduct? Again, could the answer differ depending on the product or service

⁴⁶ See especially Article 3 on access to online interfaces.

concerned? What kind of measures/tools could this be? And who should be the one providing them? A recent example from the EU lawmaking process is an idea to oblige the providers of web browsers to provide different privacy settings – as actors mediating between consumers and businesses which use tracking tools that extract information from consumers’ devices (Article 10 of the proposed ePrivacy regulation⁴⁷). One could further require businesses that apply personalized prices or display personalized ads to disclose that fact to consumers, as well as the factors that are taken into account in that context. Based on this information one could envisage a development of tools for interactive modeling that could, for example, encourage consumers to adjust their privacy settings or otherwise adjust their conduct. Yet another idea is the one of algorithmic buying groups that could be formed on the demand side to negotiate better conditions as well as other potential tools increasing consumer welfare.

Recommended readings:

Gal, M. S., & Elkin-Koren, N. (2017). Algorithmic Consumers. *Harvard Journal of Law & Technology*, 30(2), 309.

Question 8: Should we strive for an environment in which consumers always *knows* that they are being profiled, tracked, displayed targeted content, including personalized offers, and by whom? Should this information be general or specific?

If the answer to the previous two questions is negative, or only partly affirmative, would it still make sense to create technical means thanks to which consumers always knows that they are being profiled, or displayed a targeted price/commercial communication, on what basis and for what purpose? Would that be means of increasing their autonomy, or shielding them from *unknown* manipulations? What could be unintended consequences of such an environment, could it cause panic, feeling of being observed, or – on the contrary – could the ubiquity of this information cause consumers to completely ignore it? Is more information good or bad here? When and how should it be displayed to be effective (if at all)? Once again, should we rely on increasing the autonomy of consumers, i.e. more liberal approach, or just regulation of AI-deployment, i.e. a more patriarchal stand?

Recommended readings:

Mik, E. (2016). The erosion of autonomy in online consumer transactions. *Law, Innovation and Technology*, 8(1), 1.

Question 9: Should all, or at least some, automated decision-making systems be required to demonstrate procedural regularity?

In an era in which an increasing number of decisions is being taken by automated means, consumers (the civil society, oversight authorities) may want to know that such decisions are at least taken under an announced set of rules consistently applied in each case. This may be particularly true for decisions of a particular socio-economic gravity for consumers, for example related to their financial affairs. Should the algorithmic decision making systems be designed in a way to demonstrate that a given decision has not been taken on a whim, but rather resulted from a pre-defined set of conditions? If it turns out that such a requirement is unfeasible for certain types of decision-making systems (e.g. due to their constant adaptability), should businesses be prevented from using this type of systems in certain contexts?

⁴⁷ COM(2017) 10 final

Recommended readings:

Kroll, J. A., Barocas S., Felten E. W., Reidenberg J. R., Robinson D. G., & Yu H. (2017). Accountable Algorithms. *University of Pennsylvania Law Review*, 165, 633.

Question 10: Should robots/chatbots be obliged to identify themselves as such? Should they otherwise be obliged to take specific actions towards consumers, e.g. regarding purchases/recommendations made?

With an increasing number of communications between consumers and traders being conducted via artificial agents, one could consider whether there is a risk that trading parties becoming unaware of the status of their interlocutors (human or not?). If such a risk exists, the question is whether such a lack of awareness could be a cause for concern. Indeed, recent years have brought significant developments when it comes to chatbots provided by traders for purposes of communicating with consumers, as well as personal digital assistants that can handle a variety of tasks on behalf of consumers. Should such agents be obliged to identify themselves vis-à-vis their interlocutors (both consumers and traders)? As regards personal assistants deployed on the demand side: should traders be free to offer such services in any shape or form, or should there be certain baseline conditions to be followed – e.g. information duties, requirements on privacy settings / settings for the outsourcing of consumers’ tasks (making purchases and recommendations), or maybe even fiduciary duties?

Recommended readings:

Stucke, M. E., & Ezechai, A. (2017). How Digital Assistants Can Harm Our Economy, Privacy, and Democracy. *Berkeley Tech. LJ*, 32, 1239.

Brownsword, R. (2017). From Erewhon to AlphaGo: For the sake of human dignity, should we destroy the machines?. *Law, Innovation and Technology*, 9(1), 117.

Question 11: Should we have the right to “talk to a human?” In what cases or sectors?

The question, to some extent, builds upon the previous point on chatbots’ identification. Indeed, one can argue that an information requirement such as the one considered above is of little use to consumers who seek to establish contact with the trader and are not willing to speak to a machine, or are not satisfied with the outcomes of such communication. Should consumers in such cases be provided with an explicit right to “talk to a human”? Note that the EU law already provides for a fair number of information duties regarding trader’s contact details, for example the E-Commerce Directive 2000/31/EC requires that “the details of the service provider” are provided “including his electronic mail address, which allow him to be contacted rapidly and communicated with in a direct and effective manner” (Article 5(1)(c)). Is this norm effective? How would it translate to the chatbot context?

Consumers’ willingness to “talk to a human” might also go beyond their frustration about having to communicate with a machine. Perhaps more importantly, consumers are also being exposed to a growing number of decisions being taken by (fully) automated means. These range from the already discussed situations of loan denials to the blocking of user accounts on social media. Should traders be obliged to always keep a human in the loop? What for: to better understand the decisions (explanation), to reconsider the decision (appeal), to simply feel like one is not facing a blind machine (dignity)? Depending on the rationale, as well as possibly the sector and the type of consumers’ interest at stake, a different set of measures might be warranted. Some of them have been explored by the EU lawmaker in the GDPR, which the literature tends to welcome as a first step – although not a sufficient one.

Recommended reading:

Article 29 Working Party (2017a). Guidelines on Automated individual decision-making and profiling for the purposes of Regulation 2016/679.

Wachter, S., Mittelstadt, B., & Floridi L. (2017). Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76.

4.2. Technological (means)

Question 12: How do we “filter out” certain types of data from the AI-powered system’s decision-making process, if information can also be conferred indirectly? Do we intervene in data sets/ learning algorithms/ modify the outputs?

Assuming that an answer to some of the normative questions is affirmative (yes, prohibit price discrimination, discriminatory/aggressive advertising), **how does one do that?** From the machine learning point of view, the tasks realized by the systems operating the targeted commercial practices are the same: maximize the chance of a certain action (opening an ad, making a purchase) by matching a particular content (ad, price level) with a particular individual, based on data you have about behavior patterns of consumers in general (profiles), and information about a particular individual. If one wants to put constraints on the ways in which these tasks are carried out by AI-powered systems, these constraints will need to be realized, in one way or another, in the design of these systems. What are the technical options here? Is it best to intervene on the level of training data sets? Or the training algorithms? Or put some “correctors” on top of the results? Or in any other way? Is it possible to filter out certain types of information, given that with big data a lot of information might be indirectly communicated by other information (e.g. age, gender, religion or sexual orientation based on the websites on visits, or books one looks for)? What exactly are the trade-offs here? How much less efficient would the systems become as a result of this type of interventions?

Question 13: What are the useful ways of formalizing legal norms? What is the information that engineers need to do that?

Consider a simple example, based on GDPR’s prohibition of processing personal data about the individuals’ political opinions. Assume that we want AI-powered systems not to collect/utilize information about political opinions for purposes of knowledge generation or targeted communication. For a human, the term “political opinion” has a clear core and a fuzzy border. For example, “Paolo supports Party X” is a piece of information about Paolo’s political opinion, while “Paolo supports Fiorentina” probably not. However, would a statement “Pablo thinks that taxes in Italy are too low” be a political opinion? Or, “Paolo lives in the center of Rome, but secretly supports Lazio”? In all these cases, a human normative agent would be able to take into account numerous factors, without necessarily even being able to list them or make them explicit, in order to *argue* for one position or another, while at the same time often subconsciously “feeling” what an answer is in the given context. However, a machine processing data, if its command is “classify information about individuals, but when you encounter a piece of information about a political belief, disregard it”, does not have any second thoughts about the task. How to provide AI-powered systems with an equivalent of that feeling?

Question 14: How do we monitor the usage of the techniques from the previous two questions, if the outcomes are only “visible” to individual users on their individual devices? How does one oversee that which cannot be seen? A “stealth problem”

Assume a normative standard is clear – for example: disregard data about gender when displaying job advertisements; do not base ads or sexual orientation unless the consumer explicitly opts-in; never personalize diet-supplements ads using medical and/all psychological data etc. – and assume there are technical means to meet these standards – how could one *technically* check whether businesses actually comply? Unlike commercial practices on billboards, TV, in press or radio, targeted ads are visible only to a particular consumer, on their device. In case of discrimination, they might probably never notice. In case of aggressive advertisements, they might not realize. Hence, there is a need for oversight. But how does one monitor that? What are the options? Experiments (“tinkering”?), analysis of the code? How does one notice the infringement, if the infringement is hidden?

Recommended readings:

Perel, M., & Elkin-Koren, N. (2017). Black Box Tinkering: Beyond Disclosure in Algorithmic Enforcement. *Florida Law Review*, 69, 181.

Kim, P. T. (2017a). Auditing Algorithms for Discrimination. *University of Pennsylvania Law Review Online*, 166(1), 189.

4.3. Governance (means)

Finally, assuming we know what the concrete problem is, we have established what the desired state of affairs is (the goal), and we know how to technically achieve this goal, how do we design a governance system (including liability rules, rules on certification, oversight system etc.) that will convince business to comply?

Question 15: Could transparency obligations regarding targeted/personalized advertisements and offers be derived from the existing EU legal framework? If not, should they be introduced?

Building on Question 8, one cannot fail to note that the existing EU law already provides for a wide range of transparency obligations. Some of them may be found in the data protection framework (GDPR), other in consumer and marketing law *stricto sensu*. When it comes to personalized advertisements, the most directly relevant instrument appears to be the UCPD – and particularly its provisions on misleading actions and omissions. For example, Article 7(4)(c) lists the price to be paid by the consumer among the material information, which should generally be provided by the trader in the invitation to purchase. Where the nature of the product means that the price cannot reasonably be calculated in advance, at least the manner in which the price is calculated should be disclosed. None of this, however, appears to imply a disclosure duty related to the use of personalized pricing. Could such a duty be derived from a general provision of Article 7(1)? Or perhaps a respective information duty could be derived from the other parts of consumer protection *acquis* (and come within the remit of the UCPD’s misleading omissions regime via Article 7(5))? A possible candidate for such a duty is Article 6 of the Directive 2011/83/EU on consumer rights, which requires traders to disclose a list of information – e.g. about the characteristics and the total price of the goods or services and the functionality of digital content – before the conclusion of a distance contract. However, the respective provision on price is formulated in a similar way as in the UCPD and therefore seems to be subject to similar limitations. Interestingly, recital 19 of the Directive points out that “the notion of functionality should refer to the ways in which digital content can be used, for instance for the tracking of consumer behavior”. Interpreted in this way, Article 6(1)(r) of the CRD could be considered as a source of, at least, a general information duty related to targeting – somewhat similar to the one derived from the data protection framework. Nonetheless, besides a further set of questions (beginning with: when exactly is

a business-to-consumer contract for the supply of digital content concluded?), the degree to which this provision could be used to improve transparency of specific personalized offers appears to be quite limited. One could then perhaps consider rethinking the construction of the information duties in the CRD, either by redrafting the specific information items or – according to a more far-reaching proposal – by exploring the idea of the so-called “personalized disclosure”. Yet another unorthodox idea could be to treat the unilaterally drafted algorithms as standard contract terms and, as such, consider them to fall under Directive 93/13/EEC on unfair terms in consumer contracts and the price-related transparency obligations found in Article 4(2) of that act.

Against this background, do specific provisions or the broader rationale of the existing EU law support the existence of transparency obligations related to personalization and tracking mechanisms? Could such provisions be derived from the existing law and if so, are they able to produce the desired outcomes? Or would it be necessary to introduce legislative amendments? If so, should we strive for targeted amendments of the existing laws or rather look for innovative ideas?

Recommended readings:

Hacker, P. (2017). Personalizing EU Private Law: From Disclosures to Nudges and Mandates. *European Review of Private Law*, 25(3), 651.

Porat, A., & Strahilevitz, L. J. (2014). Personalizing Default Rules and Disclosure with Big Data. *Michigan Law Review*, 112(8), 1417.

Question 16: Do we need to re-think the legal framework of aggressive commercial practices? Should we prohibit them as a general rule, or bottom-up based on types of goods/ particular types of data/ cognitive biases?

If one believes – in line with Question 5 – that there is indeed a space for intervention when targeted advertisements and offers seek to exploit consumers’ vulnerabilities, what needs to be considered is how a relevant prohibition could be implemented. Should it be captured by a general rule – as it is currently in the UCPD – or rather (or perhaps additionally) be defined on a case-by-case basis? Should some of the most invasive practices be blacklisted? Should certain types of products/services be controlled more than others? Or should it rather be the types of cognitive biases that should be the criterion of prohibition? In short, how to operationalize all of this?

Recommended readings:

Helberger, N. (2016). Profiling and Targeting Consumers in the Internet of Things – A New Challenge for Consumer Law. In: R. Schulze & D. Staudenmayer (Eds.), *Digital Revolution: Challenges for Contract Law in Practice*. Nomos.

Question 17: Do we need to rethink the legal framework of non-discrimination in access to goods and services?

Even though the horizontal dimension of the EU anti-discrimination law is comparably broad, its provisions on “access to goods and services” beyond specific sectors are of limited help to the subjects of discriminatory decisions reached by algorithms. This is particularly because non-discrimination in consumer relations seems to be linked to the “general conditions of access” to a service, made available to the public at large by the provider, which apply in the absence of an individually negotiated agreement between the trader and the customer. Such a concept appears to be difficult to reconcile with the growing degrees of personalization of both commercial messages about products or services and the products and services themselves. Nevertheless, it still plays an important role in the European legislative initiatives, including the most recent ones. For example, the geo-blocking regulation 2018/302 will, as

of December 3rd 2018, apply to a situation where there are different general terms addressed at different national markets. It will not address the scenario in which prices are *personalized* depending on (among others) where the consumer comes from. Is it possible that, at some stage, there will be no “general conditions” (for some products/services)? Should we only then start to be worried about algorithmic price coordination or consider unorthodox solutions like personalized price caps? Until when is transparency or a ban on (traditional) geo-blocking enough?

Further points of criticism can be raised against the European legal framework of non-discrimination. One of such concerns is that, under existing rules, considerations of algorithmic accuracy provide an easy justification that businesses can put forward to contest allegations of indirect discrimination. The difficulties faced by potential victims in holding the owners of algorithms accountable can be further exacerbated by enforcement deficits such as the lack of access to data and algorithmic models.

A further practice affecting all sectors are the prices adjusted to the individual willingness to pay, as already, to some extent, addressed before. Does this practice give rise to additional discrimination-based concerns and, if so, is it covered by the existing non-discrimination law? Admittedly, economic factors are generally not treated as protected characteristics. But what if algorithms establish consumers’ willingness to pay based on criteria such as nationality – does this amount to direct discrimination? What if such a protected characteristics is only one of many factors? What if proxies are used? How to find out that there is a broader impact, etc.? Could it be that first-degree price discrimination is inherently non-discriminatory, considering that it goes against its “collective spirit”?

Recommended readings:

Hacker, P. (2018). Teaching Fairness to Artificial Intelligence: Existing and Novel Strategies Against Algorithmic Discrimination Under EU Law. *Common Market Law Review*, Forthcoming

Question 18: Do we need Special Data Protection Regulations? Should we govern the types of data collected, or the purposes for which they are used?

If artificial intelligence, as of now, means mostly machine learning, then regulation of “artificial intelligence” can be either regulation of algorithms, or regulation of data. The former is still ahead of us, the latter already here, in form of the GDPR – *General Data Protection Regulation*. The GDPR put in place numerous rules and principles to which data controllers and processors must adhere for the processing to be lawful. These rules and principles apply across all the sectors, all types of business models. Admittedly, a more specific set of rules applies to electronic communications (the ePrivacy Directive, currently under review), many of which, however, still appear to have the traditional telecommunication sector in mind. At the same time, arguably, processing of data generated by smart grids poses different challenges than data about shopping history, websites visited, or questions asked to digital personal assistants. Do any of these spheres require *special* data protection regulations, as for example already exist regarding medical data or telecoms?

Question 19: How to design a governance system for spheres where minor-but-widespread, stealth infringements might occur? Do we need sector/problem-specific / new governance responses?

We present this as a “single” question, but in reality it needs to be asked separately for each and every normative + technical problem. For example, if consumers see an ad based on prohibited sensitive data, might: a) not even know it is illegal; b) realize it is illegal, but fail to undertake any action, since the cost of doing something is much higher than cost of ignoring. Hence, it might be that for each and every individual consumer the harm is small, but on the societal level, the most basic principles we agreed upon as a community are being bluntly disregarded. What are the options here? Oblige businesses to conduct self-assessment drawing from the idea of accountability found in GDPR? Grant individuals

rights to compensation, allow collective actions? Rely on administrative fines? Rely on ex post (auditing / punishing infringement) or ex ante instruments (e.g. standards on how algorithms should be designed)?

Question 20: How to conceptualize the relationship between data protection law, consumer law and competition law – both at the substantive and procedural level? Do we need a digital unit in the relevant agencies so that they can cooperate better?

There is no doubt that all three domains (and likely further ones) are seriously affected by the growing importance of algorithms and AI in the market. What are the relevant overlaps and are they helpful or problematic? To what extent could these systems complement each other? How do the relevant enforcement frameworks interrelate? Is there a need for more coordination?

Recommended readings:

Costa-Cabral, F., & Lynskey, O. (2017). Family ties: the intersection between data protection and competition in EU Law. *Common Market Law Review*, 54(1), 11.

Helberger, N., Zuiderveen, B. F., & Reyna, A. (2017). The Perfect Match? A Closer Look at the Relationship between EU Consumer Law and Data Protection Law. *Common Market Law Review*, 54(5), 1427.

Kerber, W. (2016). Digital markets, data, and privacy: competition law, consumer law and data protection. *Journal of Intellectual Property Law & Practice*, 11(11), 856.

4.4. Theoretical/conceptual

At the very, very end, a couple of questions that did not fit into any other categories. They are much less “concrete” than those above. We could not refrain ourselves from posing them, though. These questions escape the regulatory logic of “what should we do” and take us rather in the direction of “what does all this mean for law and society?”.

Question 21: What could we gain from conducting research into “AI & Law & Literature”, studying cultural representations of robots and autonomous systems? Do works of (popular) art influence scholarly discourse? Do artists’ ideas give us some, even indirect, guidelines on how to proceed?

Our ideas on what artificial intelligence is often come from works of (popular) culture. It is primarily the works of fiction where we “meet” intelligent machines, observe their interactions with human beings and learn about potential opportunities and challenges of co-inhabiting the world with them. Watching *Westworld* or *I, Robot* might lead us to questions about machines’ rights. *Terminator* and *Matrix* force us to think about the existential threat to humanity that making the machines too powerful could pose. However, these works, and the images of AI presented there, appear to have little to do with what AI is today and will be tomorrow. Have these (false) images already crawled into the current legal discourses and political strategies? How to remedy this? On the other hand, even if artists portray AI differently than the reality does, is there something we could, as a society, learn from these works? Are certain questions raised by art, deep down, the questions we are facing now or will face tomorrow? Should we read the novels about robots, the market analysis, and our laws together?

Question 22: What does it mean to be a “person” in the age of “personalization”? Who/what is the person in “personalized advertising/content”? An actual consumer or his alter ego?

Personalism is one of the foundations of the Western legal culture. A “person” is at the center of the system of private law. The individual, not the collective, is the bearer of rights and duties. The distinction

between persons and things fundamentally informs the architecture of the legal system. Everyone is equal in the eyes of law, everyone is granted freedom and autonomy to lead his or her life. Hence, “personalization” of products, offers, and even rules, could seem a promising development. However, that is not what “personalization” means. With big data and AI, targeted communications and profiling, we stop being equal. Everyone gets something else. Different offers, different products, different commercials. “Personalization” means “differentiation”. Based on profiles. On data. On probability. Not on interactions with a human being. With an actual person. What does this tell us about this fundamental concept?

Question 23: How do we reconcile the conflict between freedom and autonomy on the one hand, and efficiency and convenience on the other?

20th century witnessed the heated debate between proponents of market-based economy and central planning. The question of which approach leads to more efficient social orderings have been answered by the history itself. None of the countries which attempted to create a centrally planned economy succeeded. Hence, both the deontological arguments (human dignity, freedom) and the consequentialist ones (efficiency) ended up offering support for a property & contract based social orderings favoring individual autonomy, freedom and market over planning. Probably the most famous explanation has been offered already by Hayek, who argued that knowledge about the resources available and the needs is dispersed, local and cannot be aggregated. This, however, changes with big data, Internet of Things and artificial intelligence. Currently, consumers make their preferences explicit all the time, by simply using smartphones, online platforms and services. Soon AI-powered and big data-fueled systems might become better at suggesting choices, making choices, and ultimately replacing individuals and collectives at taking economic decisions. For many, such a situation might seem convenient. Hence, as a society, we might be approaching a choice between freedom and efficiency. How to address this question?

Question 24: The vanishing line between advertising and contract, and what it implies. What does it mean for a legal system if it is difficult to tell the difference between what is still advertising and what is an offer and when the issues of personalization at pre-contractual and contractual stages become closely convergent?

Traditionally, a contractual offer and an advertisement have been treated by the legal system as two distinctive phenomena. As a result, they have been regulated by different legal instruments, informed by different principles. The first had a firmly established normative position in the contract law sources, the latter has become an object of regulation with the rise of the consumer society. Currently, in the era of AI-powered personalization, the distinction between the two is blurring. Do we need to redefine them? Or do we need a completely new conceptual framework to make sense of the existing market practice?

Question 25: Are there areas where artificial agents should behave in accordance with other (higher?) standards than humans while performing seemingly analogous tasks? For example, in the self-driving cars context, should we strive for one right ethical standard?

Artificial intelligence-powered systems perform more and more tasks traditionally undertaken by humans. These tasks are not legally neutral, i.e. they could lead to law infringement. However, the laws of today have been written with human beings as addressees in mind. And machines are not humans. In certain spheres, they can act significantly faster and be more efficient. In others, they are still far behind the capabilities of an average human. They are different, for better and for worse. The “trolley problem”, discussed in connection with self-driving cars, is a good example. We do not tell human drivers whom to sacrifice in the circumstances of an extreme and immediate decision. That is because humans act based on instinct, reflex, and cannot be expected to reason in the span of milliseconds, or to act in

accordance with pre-programmed set of rules. Machines, most probably, can. On the other hand, humans are quite good at abstract reasoning, based on general rules. “You should not process data about political opinions” is a norm that in the majority of cases will be clear and easy to implement for a human. For a machine the term “political opinion” might need to be specified to a much higher degree. One cannot simply assume the background knowledge and intuitions that humans have as a result of merely being a part of a given society and living in a particular culture. It might be that, to achieve the same goals, we need different norms for machines and for humans doing the same things. Or, it might be that our goals are different. How to go about solving this puzzle?

List of legal acts

- Council Directive 85/374/EEC of 25 July 1985 on the approximation of the laws, regulations and administrative provisions of the Member States concerning liability for defective products [1985] OJ L210/29.
- Council Directive 85/577/EEC of 20 December 1985 to protect the consumer in respect of contracts negotiated away from business premises [1985] OJ L372/31.
- Council Directive 92/59/EEC of 29 June 1992 on general product safety [1992] OJ L228/24.
- Council Directive 93/13/EEC of 5 April 1993 on unfair terms in consumer contracts [1993] OJ L95/29.
- Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data [1995] OJ L281/31.
- Directive 1999/44/EC of the European Parliament and of the Council of 25 May 1999 on certain aspects of the sale of consumer goods and associated guarantees [1999] OJ L171/12.
- Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000 on certain legal aspects of information society services, in particular electronic commerce, in the Internal Market ('Directive on electronic commerce') [2000] OJ L178/1.
- Council Directive 2000/43/EC of 29 June 2000 implementing the principle of equal treatment between persons irrespective of racial or ethnic origin [2000] OJ L180/22.
- Council Directive 2000/78/EC of 27 November 2000 establishing a general framework for equal treatment in employment and occupation [2000] OJ L303/16.
- Directive 2001/95/EC of the European Parliament and of the Council of 3 December 2001 on general product safety [2002] OJ L11/4.
- Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) [2002] OJ L201/37.
- Directive 2004/38/EC of the European Parliament and of the Council of 29 April 2004 on the right of citizens of the Union and their family members to move and reside freely within the territory of the Member States amending Regulation (EEC) No 1612/68 and repealing Directives 64/221/EEC, 68/360/EEC, 72/194/EEC, 73/148/EEC, 75/34/EEC, 75/35/EEC, 90/364/EEC, 90/365/EEC and 93/96/EEC [2004] OJ L158/77.
- Directive 2004/113/EC of 13 December 2004 implementing the principle of equal treatment between men and women in the access to and supply of goods and services [2004] OJ L373/37.
- Directive 2005/29/EC of the European Parliament and of the Council of 11 May 2005 concerning unfair business-to-consumer commercial practices in the internal market and amending Council Directive 84/450/EEC, Directives 97/7/EC, 98/27/EC and 2002/65/EC of the European Parliament and of the Council and Regulation (EC) No 2006/2004 of the European Parliament and of the Council ('Unfair Commercial Practices Directive') [2005] OJ L149/22.
- Directive 2006/54/EC of the European Parliament and of the Council of 5 July 2006 on the implementation of the principle of equal opportunities and equal treatment of men and women in matters of employment and occupation (recast) [2006] OJ L204/23.
- Directive 2006/123/EC of the European Parliament and of the Council of 12 December 2006 on services in the internal market [2006] OJ L376/36.

- Directive 2011/83/EU of the European Parliament and of the Council of 25 October 2011 on consumer rights, amending Council Directive 93/13/EEC and Directive 1999/44/EC of the European Parliament and of the Council and repealing Council Directive 85/577/EEC and Directive 97/7/EC of the European Parliament and of the Council Text with EEA relevance [2011] OJ L304/64.
- Treaty on the functioning of the European Union (consolidated version) [2012] OJ C326/47.
- Directive 2014/54/EU of the European Parliament and of the Council of 16 April 2014 on measures facilitating the exercise of rights conferred on workers in the context of freedom of movement for workers [2014] OJ L128/8.
- Directive 2014/92/EU of the European Parliament and of the Council of 23 July 2014 on the comparability of fees related to payment accounts, payment account switching and access to payment accounts with basic features [2014] OJ L257/214.
- Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) [2016] OJ L119/1.
- Regulation (EU) 2018/302 of the European Parliament and of the Council of 28 February 2018 on addressing unjustified geo-blocking and other forms of discrimination based on customers' nationality, place of residence or place of establishment within the internal market and amending Regulations (EC) No 2006/2004 and (EU) 2017/2394 and Directive 2009/22/EC [2018] OJ L 60L/1.

Bibliography

- Abbott, R. (2016). I Think, Therefore I Invent: Creative Computers and the Future of Patent Law. *Boston College Law Review*, 57, 1079.
- Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3), 488.
- Allen, T., & Widdison, R. (1996). Can computers make contracts?. *Harvard Journal of Law & Technology*, 9, 25.
- Alpaydin, E. (2016). *Machine Learning: The New AI*. MIT Press
- André, Q., Carmon, Z., Wertenbrock, K., Crum, A., Frank, D., Goldstein, W., ... & Yang, H. (2018). Consumer Choice and Autonomy in the Age of Artificial Intelligence and Big Data. *Customer Needs and Solutions*, 5(1–2), 28.
- Anrig, B., Browne, W., & Gasson M. (2008). The Role of Algorithms in Profiling. In: M. Hildebrandt & S. Gutwirth (Eds.), *Profiling the European Citizen. Cross-Disciplinary Perspectives*. Dordrecht: Springer.
- Article 29 Working Party. (2013). Opinion 03/2013 on purpose limitation. Available at: http://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2013/wp203_en.pdf
- Article 29 Working Party. (2017a). Guidelines on Automated individual decision-making and profiling for the purposes of Regulation 2016/679. Available at: http://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=612053
- Article 29 Working Party. (2017b). Guidelines on consent under Regulation 2016/679. Available at: http://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=623051
- Asaro, P. M. (2016,). The liability problem for autonomous artificial agents. *Ethical and Moral Considerations in Non-Human Agents*, 2016 AAAI Spring Symposium Series.
- Ashley, K. D. (2017). *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Cambridge University Press.
- Baldwin, R. (2014). From Regulation to Behavior Change: Giving Nudge the Third Degree. *Modern Law Review*, 77(6), 831.
- Balkin, J. M. (2015a). The Path of Robotics Law, *California Law Review*, 6, 45.
- Balkin, J. M. (2015b). Information Fiduciaries and the First Amendment. *UC Davis Law Review*, 49, 1183.
- Bar-Gill, O. (2018). Algorithmic Price Discrimination: When Demand is a Function of Both Preferences and (Mis)perceptions. *University of Chicago Law Review*, Forthcoming.
- Barnett, J., & Treleaven, P. (2018). Algorithmic Dispute Resolution. *The Computer Journal*, 61(3), 399.
- Barocas, S., & Selbst, A. D. (2016). Big Data's Disparate Impact. *California Law Review*, 104, 671.
- Barrett, L. (2017). Reasonably Suspicious Algorithms: Predictive Policing at the United States Border. *N.Y.U. Review of Law & Social Change*, 41(3), 327.
- Ben-Shahar, O., & Schneider, C. E. (2014). *More than you wanted to know: The Failure of Mandated Disclosure*. Princeton University Press.

- Bench-Capon, T., Araszkiwicz, M., Ashley, K., Atkinson, K., Bex, F., Borges, F., ... & Gordon, T. F. (2012). A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law. *Artificial Intelligence and Law*, 20(3), 215.
- BEUC. (2018). Automated Decision Making and Artificial Intelligence - A Consumer Perspective. BEUC Position Paper. Available at:
 - http://www.beuc.eu/publications/beuc-x-2018-058_automated_decision_making_and_artificial_intelligence.pdf
- Bradford, A. (2012). The Brussels Effect. *Northwestern University Law Review*, 107, 1.
- Brauneis, R., & Goodman, E. P. (2018). Algorithmic Transparency for the Smart City. *Yale Journal of Law & Technology*, 20, 103.
- Brownsword, R. (2017). From Erehwon to AlphaGo: For the sake of human dignity, should we destroy the machines?. *Law, Innovation and Technology*, 9(1), 117.
- Calo, R. (2014). Digital Market Manipulation. *George Washington Law Review*, 82, 995.
- Calo, R. (2015). Robotics and the Lessons of Cyberlaw. *California Law Review*, 103, 513.
- Caplovitz, D. (1967). *The Poor Pay More*. New York: The Free Press.
- Capoccia, G. (2015). Critical junctures and institutional change. *Advances in comparative-historical analysis*, 147-179.
- Chander, A. (2017). The Racist Algorithm?. *Michigan Law Review*, 115(6), 1023.
- Chopra, S., & White, L. F. (2011). *A legal theory for autonomous artificial agents*. University of Michigan Press.
- Citron, D. K., & Pasquale, F. (2014). The Scored Society: Due Process For Automated Predictions. *Washington Law Review*, 89, 1.
- Coase, R. H. (1960). The Problem of Social Cost. *Journal of Law and Economics*, 3, 1.
- Comparato, G., & Micklitz, H. W. (2013). Regulated Autonomy between Market Freedoms and Fundamental Rights in the Case Law of the CJEU. In: U. Bernitz, X. Groussot & F. Schulyok (Eds.), *General Principles of EU Law and European Private Law*. Aldershot: Ashgate.
- Comparato, G., Micklitz, H. W., & Svetiev, Y. (2016). The regulatory character of European private law. In: C. Twigg-Flesner (Ed.), *Research Handbook on EU Consumer and Contract Law*. Edward Elgar Publishing.
- Costa-Cabral, F., & Lynskey, O. (2017). Family ties: the intersection between data protection and competition in EU Law. *Common Market Law Review*, 54(1), 11.
- Cristianini, N. (2016). The road to artificial intelligence: A case of data over theory. *New Scientist*. Available at: <https://www.newscientist.com/article/mg23230971-200-the-irresistible-rise-of-artificial-intelligence/>
- De Búrca, G. (2012). The Trajectories of European and American Antidiscrimination Law. *The American Journal of Comparative Law*, 60, 1.
- Desai, D. R. (2015). Exploration and Exploitation: An Essay on (Machine) Learning, Algorithms, and Information Provision. *Loyola University Chicago Law Journal*, 47, 541.
- Domingos, P. (2015). *The Master Algorithm*. Basic Books.
- Doshi-Velez, F., & Kortz, M. (2017). Accountability of AI Under the Law: The Role of Explanation. Berkman Klein Center Working Group on Explanation and the Law. Available at: <https://arxiv.org/abs/1711.01134>
- Draper, N. A., & Turow, J. (2017). Audience Constructions, Reputations, and Emerging Media Technologies. In: R. Brownsword, E. Scotford & K. Yeung (Eds.), *The Oxford Handbook of Law, Regulation and Technology*. Oxford University Press.

- Dressel, J., & Farid, H. (2018). The accuracy, fairness, and limits of predicting recidivism. *Science Advances*, 4, 1.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness Through Awareness. *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 214.
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- European Data Protection Supervisor. (2015). Opinion No. 7/2015. Meeting the challenges of big data: A call for transparency, user control, data protection by design and accountability. Available at: https://edps.europa.eu/data-protection/our-work/publications/opinions/meeting-challenges-big-data_en
- European Union Agency for Fundamental Rights and Council of Europe. (2018). *Handbook on European non-discrimination law*. Available at: <http://fra.europa.eu/en/publication/2018/handbook-european-law-non-discrimination>
- Ezrahi, A., & Stucke, M. E. (2016). *Virtual Competition: The Promise And Perils Of The Algorithm-Driven Economy*. Harvard University Press.
- Ezrahi, A., & Stucke, M. E. (2017). *Artificial Intelligence & Collusion: When Computers Inhibit Competition*. Oxford Legal Studies Research Paper No. 18/2015.
- Feldman, M., Friedler, S. A., Moeller J., Scheidegger, C., & Venkatasubramanian, S. (2015). Certifying and removing disparate impact. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: how susceptible are jobs to computerisation?. *Technological forecasting and social change*, 114, 254-280.
- Fudenberg, D. J., & Villas-Boas, M. (2012). Price Discrimination in the Digital Economy. In: M. Peitz, J. Waldfogel (Eds.), *The Oxford Handbook of the Digital Economy*. Oxford University Press.
- Gal, M. S., & Elkin-Koren, N. (2017). Algorithmic Consumers. *Harvard Journal of Law & Technology*, 30(2), 309.
- Grundmann, S. (2002). Information, party autonomy and economic agents in European contract law. *Common Market Law Review*, 39, 269.
- Grundmann, S. (2016). Targeted Consumer Protection. In: D. Leczykiewicz & S. Weatherill (Eds.), *The Images of the Consumer in EU Law: Legislation, Free Movement and Competition Law*. Hart Publishing.
- Grundmann, S., Kerber, W., & Weatherill, S. (2001). Party Autonomy and the Role of Information – an Overview. In: S. Grundmann, W. Kerber & S. Weatherill (Eds.), *Party Autonomy and the Role of Information in the Internal Market*. de Gruyter.
- Hacker, P. (2017). Personalizing EU Private Law: From Disclosures to Nudges and Mandates. *European Review of Private Law*, 25(3), 651.
- Hacker, P. (2018). Teaching Fairness to Artificial Intelligence: Existing and Novel Strategies Against Algorithmic Discrimination Under EU Law. *Common Market Law Review*, Forthcoming.
- Hallevy, G. (2013). *When Robots Kill: Artificial Intelligence under Criminal Law*. UPNE.
- Haugeland, J. (1985). *Artificial intelligence: The very idea*. MIT Press.
- Hayek, F. (1945). The Use of Knowledge in Society. *The American Economic Review*, 35(4), 519.
- Helberger, N. (2016). Profiling and Targeting Consumers in the Internet of Things – A New Challenge for Consumer Law. In: R. Schulze & D. Staudenmayer (Eds.), *Digital Revolution: Challenges for Contract Law in Practice*. Nomos.

- Helberger, N., Zuiderveen, B. F., & Reyna, A. (2017). The Perfect Match? A Closer Look at the Relationship between EU Consumer Law and Data Protection Law. *Common Market Law Review*, 54(5), 1427.
- Helveston, M. N. (2016). Consumer Protection in the Age of Big Data. *Washington University Law Review*, 93(4), 859.
- Hildebrandt, M., & Gutwirth, S. (2008). *Profiling the European Citizen. Cross-Disciplinary Perspectives*. Dordrecht: Springer.
- Howells, G. (2005). The potential and limits of consumer empowerment by information. *Journal of Law and Society*, 32(3), 349.
- International Telecommunication Union. (2005). *The Internet of Things*. ITU Internet Reports 2005. Available at: http://www.itu.int/osg/spu/publications/internetofthings/InternetofThings_summary.pdf
- Janka, S. F., Uhsler, S. B. (2018). Antitrust 4.0 - the rise of Artificial Intelligence and emerging challenges to antitrust law. *European Competition Law Review*, Issue 39(3), 112.
- Jasanoff, S., & Kim, S. H. (Eds.). (2015). *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power*. University of Chicago Press.
- Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., ... & Wang, Y. (2017). Artificial intelligence in healthcare: past, present and future. *Stroke and vascular neurology*, 2(4), 230.
- Kamp, M., Körffer, B., & Meints, M. (2008). Profiling of Customers and Consumers-Customer Loyalty Programmes and Scoring Practices. In: M. Hildebrandt & S. Gutwirth, *Profiling the European Citizen. Cross-Disciplinary Perspectives*. Dordrecht: Springer.
- Kennedy, D. (2000). From the Will Theory to the Principle of Private Autonomy: Lon Fuller's "Consideration and Form". *Columbia Law Review*, 100, 94.
- Kennedy, D. (2006). Three Globalizations of Law and Legal Thought: 1850-2000. In: D. M. Trubek & A. Santos (Eds.), *The New Law and Economic Development: A Critical Appraisal*. Cambridge University Press.
- Kerber, W. (2016). Digital markets, data, and privacy: competition law, consumer law and data protection. *Journal of Intellectual Property Law & Practice*, 11(11), 856.
- Kim, P. T. (2017a). Auditing Algorithms for Discrimination. *University of Pennsylvania Law Review Online*, 166(1), 189.
- Kim, P. T. (2017b). Data-Driven Discrimination at Work. *William & Mary Law Review*, 48(3), 857.
- King, N. J., & Forder, J. (2016). Data analytics and consumer profiling: finding appropriate privacy principles for discovered data. *Computer Law & Security Review*, 32(5), 696.
- Kochenov, D., de Búrca, G., & Williams, A. (Eds.) (2015). *Europe's Justice Deficit?* Oxford: Hart Publishing.
- Krishnan, A. (2016). *Killer robots: legality and ethicality of autonomous weapons*. Routledge.
- Kroll, J. A., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2017). *Accountable Algorithms*. *University of Pennsylvania Law Review*, 165, 633.
- Kullmann, M. (2018). Platform Work, Algorithmic Decision-Making, and EU Gender Equality Law. *International Journal of Comparative Labour Law and Industrial Relations*, 34(1), 1.
- Kurzweil, R., Richter, R., Kurzweil, R., & Schneider, M. L. (1990). *The age of intelligent machines* (Vol. 579). Cambridge, MA: MIT press. .
- Langely, P. (2014). Equipping entrepreneurs: consuming credit and credit scores. *Consumption Markets & Culture*, 17(5), 448.

- Larson, D. A. (2010). Artificial Intelligence: Robots, Avatars, and the Demise of the Human Mediator. *Ohio State Journal on Dispute Resolution*, 25(1), 105.
- Levine, M. E. (2002). Price Discrimination Without Market Power. *Yale Journal on Regulation*, 19(1), 1.
- Lightbourne, J. (2017). Algorithms & Fiduciaries: Existing and Proposed Regulatory Approaches to Artificially Intelligent Financial Planners. *Duke Law Journal*, 67(3), 651.
- Lindsay, A., & McCarthy, E. (2017). Do we need to prevent pricing algorithms cooking up markets?. *European Competition Law Review*, 38(12), 533.
- Lösch, A., & Schneider, C. (2016). Transforming power/knowledge apparatuses: the smart grid in the German energy transition. *Innovation: The European Journal of Social Science Research*, 29(3), 262.
- Mak, V. (2013). The “Average Consumer” of EU Law in Domestic and European Litigation. In Dorota Leczykiewicz and Stephen Weatherill (eds.), *The Involvement of EU law in Private Law Relationships*, Hart Publishing.
- McCrudden, C., & Prechal S. (2009). The Concepts of Equality and Non-Discrimination in Europe: A practical approach. *Oxford Legal Studies Research Paper No. 4/2011*. Available at: <http://ec.europa.eu/social/BlobServlet?docId=4553>
- Mehra, S. K. (2015). Antitrust and the Robo-Seller: Competition in the Time of Algorithms. *Minnesota Law Review*, 100, 1323.
- Micklitz, H. W. (2006). The general clause on unfair practices. In: G. Howells, H. W. Micklitz, & T. Wilhelmsson (Eds.), *European fair trading law: The unfair commercial practices directive*. Routledge.
- Micklitz, H. W. (2011). Social Justice and Access Justice in Private Law. *EUI Working Paper LAW 2011/02*. Available at: <http://cadmus.eui.eu/handle/1814/15706>
- Micklitz, H. W. (2012). The Expulsion of the Concept of Protection from the Consumer Law and the Return of Social Elements in the Civil Law. A Bittersweet Polemic. *Journal of Consumer Policy*, 35, 283.
- Micklitz, H. W. (2015). On the Intellectual History of Freedom of Contract and Regulation. *EUI Working Paper LAW 2015/09*. Available at: <http://cadmus.eui.eu/handle/1814/35178>
- Micklitz, H. W., & Patterson, D. (2013). From the nation state to the market: the evolution of EU private law. In: B. van Vooren, S. Blockmans & J. Wouters (Eds.), *The EU's role in global governance: the legal dimension*. Oxford University Press.
- Micklitz, H. W., & Reich, N. (2014). The Court and Sleeping Beauty: The revival of the Unfair Contract Terms Directive (UCTD). *Common Market Law Review*, 51, 771.
- Mik, E. (2016). The erosion of autonomy in online consumer transactions. *Law, Innovation and Technology*, 8(1), 1.
- Mik, E. (2017). Smart contracts: terminology, technical limitations and real world complexity. *Law, Innovation and Technology*, 9(2), 269.
- O’Neill, C. (2016). *Weapons of math destruction. How Big Data Increases Inequality and Threatens Democracy*. Broadway Books.
- Pagallo, U. (2013). *The laws of robots: crimes, contracts, and torts*. Vol. 10. Springer Science & Business Media.
- Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press.
- Pasquale, F., & Citron, D. K. (2014). Promoting Innovation While Preventing Discrimination: Policy Goals for the Scored Society. *Washington Law Review*, 89, 1413.

- Pearce, H. (2015). Online data transactions, consent, and big data: technological solutions to technological problems?. *Computer and Telecommunications Law Review*, 21(6), 149.
- Perel, M., & Elkin-Koren, N. (2017). Black Box Tinkering: Beyond Disclosure in Algorithmic Enforcement. *Florida Law Review*, 69, 181.
- Podszun, R., & Kreifels, S. (2016). Digital platforms and competition law. *Journal of European Consumer and Market Law*, 5(1), 33.
- Porat, A., & Strahilevitz, L. J. (2014). Personalizing Default Rules and Disclosure with Big Data. *Michigan Law Review*, 112(8), 1417.
- Rayna, T., Darlington J., & Striukova L. (2015). Pricing music using personal data: mutually advantageous first-degree price discrimination. *Electronic Markets*, 25(2), 139.
- Reich, N. (2013). The Impact of the Non-Discrimination Principle on Private Autonomy. In: D. Leczykiewicz, S. Weatherill (Eds.), *The Involvement of EU Law in Private Law Relationships*. Hart Publishing.
- Reich, N. (2014). *General Principles of EU Civil Law*. Intersentia.
- Rich, M. L. (2016). Machine Learning, Automated Suspicion Algorithms, and the Fourth Amendment. *University of Pennsylvania Law Review*, 164, 871.
- Rudin, C. (2014). Algorithms for Interpretable Machine Learning. Video lecture available at: http://videlectures.net/kdd2014_rudin_machine_learning/
- Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*. Pearson Education Limited.
- Sartor, G. (2009). Cognitive automata and the law: electronic contracting and the intentionality of software agents. *Artificial Intelligence and Law*, 17(4), 253.
- Sax, M., Helberger, N., & Bol, N. (2018). *Journal of Consumer Policy*, 41, 103.
- Scherer, M. U. (2015). Regulating artificial intelligence systems: Risks, challenges, competencies, and strategies. *Harvard Journal of Law & Technology*, 29, 353.
- Schmitz, A. J. (2014). Secret Consumer Scores And Segmentations: Separating “Haves” From “Have-Nots”. *Michigan State Law Review*, 1411.
- Simitis, S. (2010). Privacy – An Endless Debate?. *California Law Review*, 98, 1989.
- Steppe, R. (2017). Online price discrimination and personal data: A General Data Protection Regulation perspective. *Computer Law & Security Review*, 33, 768.
- Stigler, G. (1968). *The Theory of Price*. Collier-Macmillan.
- Stucke, M. E., & Ezechia, A. (2017). How Digital Assistants Can Harm Our Economy, Privacy, and Democracy. *Berkeley Technology Law Journal*, 32, 1239.
- Tonner, K. (2014). From the Kennedy Message to Full Harmonising Consumer Law Directives: A Retrospect. In: K. Purnhagen & P. Rott (Eds.), *Varieties of European Economic Law and Regulation: Liber amicorum for Hans Micklitz*. Springer
- Tridimas, T. (2007). *The General Principles of EU Law*. 2nd edn. Oxford University Press.
- Vedder, A., & Naudts, L. (2017). Accountability for the use of algorithms in a big data environment. *International Review of Law, Computers & Technology*, 31(2), 206.
- Valdman, M. (2010). Outsourcing self-government. *Ethics*, 120(4), 761–790.
- Vladeck D. C. (2016). Consumer protection in an era of big data analytics. *Ohio Northern University Law Review*, 42, 493.
- Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76.

- Weatherill, S. (2001). Justifying Limits to Party Autonomy in the Internal Market – EC Legislation in the Field of Consumer Protection. In: S. Grundmann, W. Kerber & S. Weatherill (Eds.), *Party Autonomy and the Role of Information in the Internal Market*. de Gruyter.
- Weatherill, S. (2013). The Elusive Character of Private Autonomy in EU Law. In: D. Leczykiewicz & S. Weatherill (Eds.), *The Involvement of EU law in Private Law Relationships*. Hart Publishing.
- Weber, R. H., & Weber, R. (2010). *Internet of Things: Legal Perspectives*. Springer.
- World Wide Web Foundation. (2017). *Algorithmic Accountability: Applying the concept to different country contexts*. Available at: https://webfoundation.org/docs/2017/07/Algorithms_Report_WF.pdf
- Worthington, S. (2015). *Common law values: The role of party autonomy in private law*. University of Cambridge Faculty of Law Research Paper No. 33/2015.
- Yeung, K. (2017). ‘Hypernudge’: Big Data as a mode of regulation by design. *Information, Communication & Society*, 20(1), 118.
- Young, R. (2017). *Personal autonomy: Beyond negative and positive liberty*. Vol. 9. Taylor & Francis.
- Zarsky, T. (2014). Understanding Discrimination in the Scored Society. *Washington Law Review*, 89, 137.
- Zarsky, T. (2016). The Trouble with Algorithmic Decisions. An Analytic Road Map to Examine Efficiency and Fairness in Automated and Opaque Decision Making. *Science, Technology, & Human Values*, 41(1), 118.

