

# *EU Consumer Law and Artificial Intelligence*

AGNIESZKA JABŁONOWSKA<sup>1</sup>  
PRZEMYSŁAW PAŁKA<sup>2</sup>

## **Abstract**

The chapter considers the interplay between the EU consumer law and artificial intelligence. It begins by addressing three topical issues in law & AI scholarship: the definition of “artificial intelligence”, the problem of “black box” and the liability for actions of AI. It addresses certain misconceptions present in legal discourses, and proposes conceptual distinctions to remedy them. The argument then moves to three issues specific to consumer law: the asymmetry between businesses and consumers, the emergence of AI applications in consumer markets, and some of the potential challenges which they raise, particularly those to consumers’ autonomy. Deeper intellectual engagement with the challenges posed to the doctrine and the practice of consumer law by the development of AI is necessary before one can move to recommendations on how the law should be changed. The chapter concludes with a plea to develop AI-based applications empowering the consumers and the civil society.

## I. Introduction

Artificial Intelligence is about to go mainstream in the European legal scholarship. Unquestionable success of this technology across different sectors of the economy – from energy, to retail, to healthcare – led to a proliferation of scholarly contributions and political strategies during the last couple of years. The next EU Commission is expected to make the AI one of its strategic priorities, akin to Juncker’s Digital Single Market. The audience interested in AI will therefore gradually broaden, from scholars interested in technology, to scholars working in particular legal fields, where AI will become one of many phenomena to be taken into account, as was the case with the Internet. The societal changes caused by AI are significant, and so research concentrating on identification of promises and challenges, as well as possible ways to respond to them, is important and should be applauded. At the same time, the tone and content of some of these contributions suggests that artificial intelligence is a hype. Hence, many predictions made about its course, as well as concerns raised, should be treated with critical reservation. In this Chapter, we want to draw from the lessons we have learned during our work in the last couple of years, and anticipate some questions and issues we believe will be discussed over the next years by European lawyers, especially those interested in consumer law and policy.

---

<sup>1</sup> Ph.D. researcher at the University of Lodz. Supported by the Foundation for Polish Science (FNP).

<sup>2</sup> Private Law Fellow at Yale Law School’s Center for Private Law, Resident Fellow at Information Society Project.

In particular, we build upon the results of the ARTSY Project<sup>3</sup> led by Hans Micklitz for the last two years, in which we had a pleasure to participate. Its ambition was to map the ways in which the AI is being used in the private sector, and identify the challenges for the EU consumer law and policy. The final results are presented in the form of a comprehensive report<sup>4</sup> and high-level takeaways,<sup>5</sup> and the Reader is kindly invited to consult them. We do not want to replicate the Project’s conclusions here. Rather, we build upon them and offer reflections about three substantive matters we consider scholarly interesting and societally important: the asymmetry between businesses and consumers, the emergence of AI applications in consumer markets as well as some of the potential challenges which they raise. These reflections are the most consumer-oriented part of the Chapter, presented in sections IV-VI.

Hans Micklitz’s contribution to the scholarly community is not limited to the impressive substantive results of his studies. Equally important is his way of thinking. The urge to be critical, and to go against the current. To be sceptical of hypes and herd behaviour. To think outside of the box, and break down the established labels. To take time to observe and reflect, before rushing to conclusions. Along these lines, we begin the Chapter by tackling three “topical” issues in the field of AI & law: the problem of its definition, the problem of the “black box” and the problem of liability for AI’s actions. To do so, we open by a provocation, and want to argue that neither of them is actually a problem. Rather, we tend to see them as problems as a result of conceptual confusions. We clarify the concepts, propose several distinctions to help structure the debates, and suggest what are the alternative ways of tackling these questions.

Finally, we want to give due respect to Hans Micklitz’s deep commitment to the society, and the belief that as scholars, we owe a debt to the people around us. In his work, Hans beautifully connects the ability to slow down and theorize, with the conviction that at the end of the day, what we do should serve a societal purpose. That action research is something not only compatible, but also deeply intellectually connected with the theory. To do so, we briefly reflect upon the ways in which artificial intelligence can be used to empower the consumers and the civil society, what are the legal and political preconditions for this to happen, and so what should we do next.

## II. Definition of Artificial Intelligence. Or, What Is This “AI”?

It seems to be a good practice in the field of AI & law to begin a paper by stating that there is no commonly agreed upon definition of AI. This statement is correct, but hardly helpful. In the end,

---

<sup>3</sup> ARTSY Project (ARTificial intelligence SYstems and consumer law & policy) was conducted over the course of 2017 and 2018 at the European University Institute in Florence, with Hans Micklitz being the Principal Investigator. It was funded by the Berkman Klein Center for Internet and Society as a part of the Ethics and Governance of Artificial Intelligence Fund. See: <https://artsy.eui.eu/>.

<sup>4</sup> A Jabłowska, M Kuziemski, AM Nowak, H-W Micklitz, P Pałka and G Sartor, ‘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’ (2018) EUI Working Paper LAW 2018/11. Available at: <http://cadmus.eui.eu/handle/1814/57484>.

<sup>5</sup> P Pałka, A Jabłowska, H-W Micklitz and G Sartor, ‘Before Machines Consume the Consumers: High-Level Takeaways from the ARTSY Project’ (2018) EUI Working Paper LAW 2018/12. Available at: <http://cadmus.eui.eu/handle/1814/57485>.

an author and the readers should share an understanding of what is it that they are talking about. The lack of *the* definition does not remove the author’s obligation to clarify what is the object of the inquiry. Given that one can expect the number of legal papers concerned with AI to grow, it is worthwhile to begin our considerations by addressing the questions: why is there no agreed upon definition of AI?; and: what are the **alternative** ways to begin a paper, which actually confer some added knowledge?

Let us begin by stating that a lack of a definition is not a problem in itself. There are quite a few fundamental concepts: “human being”, “politics”, “economy”, or – yes – “law”, which lack a commonly agreed upon definition, what does not stop us from studying them or having some real world impact upon them. What matters, however, are the reasons why scholars would like to have one. What are the functions a definition is supposed to serve? And there are two. First, the role of a definition is to delineate the object of inquiry. What is it that we are talking about? When we say “artificial intelligence”, do we mean “self-driving cars”, “a discipline in engineering”, “Google translate”? This is something we need to agree upon, to avoid talking past one another. Second, once the object of inquiry is clarified, a role of a definition is to tell us something about it. What are the features/characteristics of the phenomenon under consideration? This is something we should do, to avoid talking about our (possibly mistaken) idea about the reality, instead of the reality itself. So the reflex to define – in order to delineate and increase understanding – should be applauded. However, these functions do not have to be served by an Aristotelian definition *per genus et differentiam*. There are other ways. A lack of such a definition does not remove an obligation to clarify. It is helpful, however, to inquire into why is there no commonly agreed upon definition of AI. There are two reasons.

First, the same term is currently being used to refer to very different phenomena. And this is not a bad way to start. We should not begin research by being underinclusive, for the sake of a neat definition. The buzzword “artificial intelligence” captures a series of intuitions, fears and hopes in people’s minds. Scholars who try to generalize sometimes have very different things in mind. Some will be concerned with robots becoming conscious, like those in *I, Robot* or *Westworld* (which are not coming about anytime soon) and talk about “robot rights” or “personality for AI”. Some others will think of high-frequency trading agents, taking decisions “autonomously” (which have been with us for a decade now) and ponder the need to update financial markets regulation. There are scholars concerned with credit-scoring by algorithms (which are real, but technologically simple, and for now do not rely on AI), and there are those interested in moral considerations made salient by self-driving cars.<sup>6</sup> All of these phenomena have something in common with artificial intelligence. At the same time, they are very different from one another. Some exist already, some do not. Some pose threats to health and life, some to other legally protected values. Just like we do not define a “human intelligence” for the sake of regulating human behaviour, there is no reason to believe that we will need *the* definition of AI, either for scholarship, or for regulation.

---

<sup>6</sup> G Contissa, F Lagioia and G Sartor, ‘The Ethical Knob: Ethically-Customisable Automated Vehicles and the Law’ (2017) 25 *Artificial Intelligence and Law* 365.

Second, perspective matters. An engineer sitting in a lab and trying to build a computer that will pass the Turing test<sup>7</sup> has something else in mind than a policymaker trying to understand what are the problems with companies using AI to deliver targeted advertising to consumers. Consider three definitions, first two commonly invoked after Russell and Norvig’s classic textbook,<sup>8</sup> and the third used by us in the ARTSY project:

1. ‘The art of creating machines that perform functions that require intelligence when performed by people’;<sup>9</sup>
2. ‘The study of computations that make it possible to perceive, reason and act’;<sup>10</sup>
3. ‘A socio-technological practice where an entity, private or public, uses machine learning tools to generate new knowledge out of large amounts of data, and/or act upon that knowledge, in order to automate and optimize certain process, as well as undertake new, previously impossible, tasks’.<sup>11</sup>

The first two define an academic discipline, or a goal of a scholarly endeavour. The object of inquiry is what engineers do in their labs. The *genus* being “the art” or “the study”, clearly indicates that.<sup>12</sup> By comparing definitions 1 and 2, one can see what is the disagreement in the engineering community. The first camp is interested in creating machines that will act like humans, deliver outputs which look humanly, even if achieved through very different means. The second is concerned with processes, and tries to replicate human way of thinking inside a machine. Both of them, in addition, have a universal ambition.

The third definition is different in two ways. First, its *genus* is “practice”. It is not concerned with what engineers try to achieve, but what is actually happening on the market right now.<sup>13</sup> Second, it does not have a universal ambition. It defines the term for the purposes of studying the practice from the point of view of consumer law. It clearly delineates what is meant, at the same time being open to the idea that for the purposes of studying self-driving cars, autonomous weapons or a field of investment, “AI” will be defined differently. Differently, because a different set of phenomena will be under consideration, and because that set might have different unifying features.

The reader will have noticed the risk of defining AI *ignotum per ignotius* in the third definition. Its understanding presupposes the knowledge of what machine learning is. Why is this so, and what is it?

---

<sup>7</sup> A test for an “intelligent machine” has been proposed by Alan Turing, who argued that we could deem one to be so, if a human talking to it (by textual means) was unable to tell whether the entity on the other side is a human or a machine.

<sup>8</sup> SJ Russell and P Norvig, *Artificial Intelligence: A Modern Approach*, 3rd edn (Pearson Education Limited, 2016).

<sup>9</sup> R Kurzweil, *The Age of Intelligent Machines* (Cambridge, MA, MIT Press, 1990).

<sup>10</sup> PH Winston, *Artificial Intelligence*, 3rd edn (Addison-Wesley, 1992).

<sup>11</sup> Pałka, Jabłonowska, Micklitz and Sartor, ‘Before Machines’ (2018).

<sup>12</sup> If one wanted to define “law” in a similar manner, one could speak of “the art of commenting on legal material in a way that finds order, provides guidance, and expresses critical take on the subject”. The study of something.

<sup>13</sup> If one wanted to define “law” in a similar manner, one could speak of “the ways in which legal actors, like judges or attorneys behave, and the reasons they provide for doing so”. The actual practice.

For the purposes of pondering effects of technology on law, and considering potential regulation, one should concentrate on what is possible/existent, or what will be so in the foreseeable future. Machine learning is just one type of artificial intelligence,<sup>14</sup> but the only type that currently works and is used in the practice.<sup>15</sup> That is why it makes sense for legal scholars to be clear about the types of technologies they refer to when invoking the term “AI”. What is machine learning, then?

Machine learning puts upside-down the traditional approach to computer science.<sup>16</sup> Traditionally, a programmer would write a series of instructions for the machine, which the latter would subsequently execute. For certain tasks – like creating spreadsheets – this works well enough. But for others – like machine translation or image recognition – traditional programming came short of rendering satisfactory results. In machine learning, a developer will give a computer a big data set – a series of inputs and outputs one could say – and let the machine, relying on a learning algorithm, to figure out the way to get there. For example, to teach a computer to recognize traffic signs, an image recognition learning algorithm will receive one billion photos of such signs, and one billion photos without them. The output will be a trained algorithm, a new piece of code capable of conducting a certain task. These algorithms, later, can be incorporated in a piece of software (an app, a website, a desktop computer program) that a user can employ.

This approach has now successfully been deployed in numerous ways: machine translation, image recognition, spam detection, health diagnosis, fraud detection etc. And this is the “artificial intelligence” that “is everywhere” nowadays. Which points to the final question: how should one then begin a paper about law and AI?

First, it seems to be a good idea to specify what exactly one has in mind: current practice or a possible future? If the former, to clarify what exactly are the techniques used in a given sphere, by whom, and for what purpose. If the latter, to rely on technical research making clear whether we are talking about foreseeable applications (like new types of natural language processing), or far removed speculations (like conscious machines). For example, “this paper is about the machine learning tools advertisers use to better understand consumer reaction to ads, and in consequence develop more effective ones”. The beauty of this approach is that it gets real. The difficulty is that it requires empirical research, often interdisciplinary in nature, which is much more effort-intensive than armchair philosophy. But it is the way to go.

Second, one has to remember that AI is not a one thing, it does not have one definition, does not pose the same types of problems everywhere, and so will have to be tackled case-by-case. The times when one could speak of “regulation of AI” and come across with a meaningful contribution, are over.

What exactly do you want think of, when you want to speak of “Artificial Intelligence”?

---

<sup>14</sup> E Alpaydin, *Machine Learning. The New AI* (MIT Press, 2016).

<sup>15</sup> N Cristianini, ‘The Road to Artificial Intelligence: A Case of Data over Theory’ (2016) *The New Scientist*. Available at: <https://www.newscientist.com/article/mg23230971-200-the-irresistible-rise-of-artificial-intelligence/>.

<sup>16</sup> Alpaydin, *Machine Learning* (2016).

### III. Black Box

A commonly repeated wisdom is that programs relying on artificial intelligence, or even better – on “algorithms” – are “black boxes”. We do not know what is going on inside, how they perform their tasks. Some people express fear that we will never know. Others set up for the task of “opening” black boxes and advocate for “the FDA for algorithms” or “more research on explainability”. However, just like “AI”, “black box” is a buzzword used frequently to denote different things. In this section, we want to introduce two distinctions: between legally- and technically induced black boxes; and between private and public use of algorithms. And we argue that even though there are many challenges surrounding different algorithmic technologies, the “black box problem” is much smaller a problem (at least from regulatory perspective) that people tend to assume.

To begin, let us address yet another concept that often gets used rather frivolously: “algorithm”. An algorithm, from the technical standpoint, is a sequence of unambiguous instructions to be followed by an agent. This agent can be a human, or a computer (an artificial agent). Consider an example. Imagine your friend is hosting a party, but there is too much stuff to do, so you volunteer to help. She asks you to offer drinks to the guests. You’re stressed, and do not know how to go about it, so your friend gives you the following instruction:

1. Look at the guests, and when you see one not holding a drink, approach and ask whether he or she wants one;
2. IF NOT, exclude this guest from the pool for 20 minutes and go back to step 1.
3. IF YES, ask what drink would they like.
4. IF we have that drink, serve it, wish them a lot of fun, and go back to step 1; IF we do not have that drink, proceed to step 5
5. List the drinks which we have. IF the guest chooses one, serve it, wish them a lot of fun, and go back to step 1; IF the guest does not want one anymore, execute the step 2.

The instruction above is an algorithm. Now, the agent executing it is you – a human – but it could just as well be a robot. All this is to say than an algorithm and an agent (often an artificial agent) executing it is not the same type of entity.

A “black box” is an agent whose instructions we do not know. There are two reasons why do not know it (not mutually exclusive): because the agent, or its employer/owner does not want to tell us what the instruction is (and claims to have a legal title not to tell us), or does not understand themselves what the instruction is. Those, respectively, legal and technical black boxes.

Legal black boxes emerge when a company who developed or uses a tool invokes a trade secret, or another legal reason, to not share the instruction (the algorithm) that the tool follows. This, often, can be quite simple algorithms. In fact, many algorithms labelled “black boxes” do not rely on artificial intelligence. Consider two examples that captured significant public attention in the past: the German SCHUFA case, and the COMPAS risk-assessment tool used in the American criminal justice system. That is the “proprietary” nature defence – “I came up with this algorithm, it took me time and effort, and I will not share, because others will replicate it”. In many cases,

such a defence makes sense. In other cases, there will be public policy reasons for not disclosing the instructions. Algorithms used by agents at the airports that choose whom to control, or by tax authorities choosing whom to inspect, probably should remain secret, in order not to be “gamed” by the bad guys. However, in all this cases it is technically possible to understand the algorithm by humans. And there are many ways in which, in order to make sure that there is no discrimination, or other undesirable consequences, they could be controlled by experts, without full disclosure to the public. One can imagine administrative agencies overseeing financial markets, or consumer relations, being given a competence to inspect them, while on the same time being legally obliged not to disclose. Especially for algorithms that are privately developed, but used as a part of the exercise of public power (by the criminal justice system, tax authorities, in visa lotteries etc.) there should be someone competent (legally and technically) to audit them. To achieve this, however, no rocket science is needed; only changes in law and organization.

Technical black boxes are those that emerge as a result of a learning algorithms finding patterns in vast collections of data. And indeed, in software relying on AI, certain steps are “mysterious” to a human, in the sense that we do not really know how the machine gets from point A to point B. Many examples caught media attention. The, by now famous, AlphaGO has learned how to beat human masters in the Go game, first by studying thousands of games of amateurs and professionals,<sup>17</sup> and then got even better simply by playing against itself.<sup>18</sup> Human experts do not know **how** the machine is “thinking”, while at the same time it outperforms every human on the planet. DeepMind, AlphaGo developers, captured the general public shock on their website, writing:

During the games, AlphaGo played a handful of highly inventive winning moves, several of which – including move 37 in game two – were so surprising they overturned hundreds of years of received wisdom, and have since been examined extensively by players of all levels. In the course of winning, AlphaGo somehow taught the world completely new knowledge about perhaps the most studied and contemplated game in history.<sup>19</sup>

The game example is, well, a toy example but the similar technical black boxes occur in the systems with high societal and practical relevance. Medical applications detecting cancer, or deciding whether to use chemotherapy; self-driving cars learning how to drive and how to behave in different situations; machine translation (potentially soon to be used in courts and administrative proceedings) all have this “mystery” part in them, where we do not know how exactly they perform their tasks. Is it a problem for regulation? Much smaller than we tend to assume. Because, there is

---

<sup>17</sup> D Silver, A Huang, CJ Maddison, A Guez, L Sifre, G van den Driessche, J Schrittwieser, I Antonoglou, V Panneershelvam, M Lanctot, S Dieleman, D Grewe, J Nham, N Kalchbrenner, I Sutskever, T Lillicrap, M Leach, K Kavukcuoglu, T Graepel and D Hassabis, ‘Mastering the Game of Go with Deep Neural Networks and Tree Search’ (2016) 529 *Nature* 484.

<sup>18</sup> D Silver, J Schrittwieser, K Simonyan, I Antonoglou, A Huang, A Guez, T Hubert, L Baker, M Lai, A Bolton, Y Chen, T Lillicrap, F Hui, L Sifre, G van den Driessche, T Graepel and D Hassabis, ‘Mastering the Game of Go without Human Knowledge’ (2017) 550 *Nature* 354.

<sup>19</sup> See: <https://deepmind.com/research/alphago/>.

really a lot that we do know. We know why the machine is trying to get to point B from point A; and in what circumstances it has learned to do so, and who was in charge of these circumstances.

Imagine that a small child, who cannot read or write, points to a jar labelled “COOKIES” and asks what the sign says. You answer, and a child spends a while looking at it. On a different day, it takes a piece of paper, scribbles something, and tells “this says *COOKIES*”. Indeed, that is what the output looks like. You’re happy and proud, congratulate the child, and ask it to write “cook”. “I have no idea how to write that” is the answer. You ask the child to show you the letter “c”. “What’s a letter?” you hear in reply. The child does not know the letters, it cannot “write” in the sense of the word that you would conceive of. Nevertheless, it has learned to create a sign indicating “cookies”. You do not know *how* it learned to so. Definitely not the “standard” way. What the child paid attention to, and how it managed to remember the “cookies” shape is a mystery to you. A black box. At the same time, you know what data the child was exposed to. The sign “COOKIES”, paired with the meaning of the sign that you provided. And you know what the task was: to learn how to write “cookies”.

The same is often the case with most machine learning systems. Their programmers/operators might not know how precisely the machine has figured out how to perform the task at hand; but they know quite well what (types of) data it had access to, or what was the task it was supposed to learn to perform. Hence, it is a significant overstatement to say that “no-one understands” what the machine is doing. Of course, in many applications, it would come useful to know a little bit better how exactly the machine has learned what to do. From the perspective of the computer science, this is a widely acknowledged challenge, and a promising venue of research. However, from the regulatory point of view, this is often a third order problem. What matters for regulation is: what tasks the developers want the machines to execute; in what ways have they trained them; and, in case there is a risk of a known undesirable output (discrimination, for example), what precautions at the implementation stage have been undertaken (the algorithm is not the same thing as the agent; the agents might have additional layers of instructions, not necessarily relying on AI). To use yet another distinction and popular buzzword, what we are dealing with is first and foremost organisational opacity, not a technical one.<sup>20</sup> Instead of focusing on “explainability” of a “black box”, we should focus on the processes of development and deployment, and a vast majority of answers can be found there.

With so many different types of black boxes, again the “problem” is a legion, and will need to be tackled case by case, often using very different legal and technical means. If one wishes to write about black box in a particular societal practice, a good way to start is to ask: what type of a black box are we dealing with here; and: what exactly is the problem caused by it.

#### IV. Liability

---

<sup>20</sup> See: H Chulu, ‘Let Us End Algorithmic Discrimination’ (2018) *Medium*. Available at: <https://medium.com/techfestival-2018/let-us-end-algorithmic-discrimination-98421b1334a3>.

The third issue often invoked at the early stages of discussions about legal status and regulation of artificial intelligence is the issue of liability. A familiar story would begin by warning us that machines are getting “smarter and smarter”, “autonomous”, “reprogram themselves”, and so as a result no one knows what exactly they will do. In such a world, the narrative continues, who is/should be liable? The owner? The seller? The programmer? The question is difficult and complex, it would seem, and so this paradigmatic shift, “the responsibility gap”, needs to be addressed soon by researchers and policy. Or so we are told.

As we hope is clear by now, many of these assumptions are quite stretched. Granted, tools based using artificial intelligence are more sophisticated than traditional software, and their actions are harder to predict; but we are very far from a world in which machines get “autonomous” in the sense in which they start doing things by themselves. How they achieve their goals might be difficult to explain, but it is still very much humans who give them the goals, and are capable of foreseeing negative consequences. In this sense, an action of an AI-powered system is never an accident. It is always a result of a quite informed decision, taken by a human, to deploy the tool in a given context, with clearly specified objectives in mind. Hence, we want to argue, liability gap is much smaller a problem than many tend to assume. Nevertheless, the question is there to be answered. To do so, we wanted to introduce two distinctions: between moral and legal liability; and between compensatory and preventive function of legal liability.

As a matter of moral philosophy, the question of who is *actually* culpable if a machine relying on AI does something wrong, is interesting and arguably challenging. However, it is a very separate question from the legal one. As a matter of law, a situation in which we ascribe civil liability to an entity who had no fault is neither new nor challenging. Liability for actions of children or animals is one example. Strict liability for “dangerous” products is another. Is the producer of Coca Cola cans morally liable for the fact that one in billion exploded? Not really. Is he legally liable? Fully. In yet another domain, traffic law provides a good example. Quite often the person who was actually at fault – driving without slowing down on the main road – will not be found guilty by the law, as long as he or she had a right of way. Traffic law is a good starting point, because it shows how we managed to fully separate the moral question from the legal one, and on the same time ensure that the law’s compensatory function is realised. A system of strict liability, with a recursive claim, paired with mandatory insurance at various stages, shows that the question “who pays?” can be both answered unambiguously, and fully separated from the question of who is “actually” morally liable.

However, ensuring that people harmed are duly compensated is not the only function of legal liability. Arguably, as a society, we want to live in a world in which someone who was harmed receives compensation; but we would prefer a world in which that someone does not get harmed in the first place. Hence, liability also serves a preventive function. People do not want to be liable and have to pay; and so they act more carefully, or do not commit certain actions in a first place. With machines, however, this will not necessarily be the case. Machines need to be programmed/trained to act in a certain way or not. And creating an incentive to design them in such a way does not have to come in a form of rules on liability. Many other options are on the

table. Mandatory or optional certification, certain good practices (like impact assessments) already in the design process, transparency and accountability mechanisms etc. are just some examples of other means of achieving the goal of machines not breaking the law.

What is very important is the realization that an exact choice of means will be different in different cases, both for technical and for normative reasons. Sometimes creation of standards will be possible and helpful, sometimes not. Sometimes keeping a “human in the loop” is possible, sometimes very difficult. Add sometimes (when life and health are at risk) we might want to have these costly processes legally required, while in other cases we might think that the cost is too high to bear.

The question of liability is therefore a problem of complexity of situations, not complexity of technology. When encountering one, it is good to remember that 1) moral and legal responsibility are not necessarily connected to one another; 2) ensuring that there will be someone to pay is not equal to ensuring that law’s preventive function is realised; and 3) liability rules (and other preventive measures) can and will differ case by case, depending on technical and normative divergences.

To sum up the argument until now: the three issues discussed above, namely the definition of Artificial Intelligence, the problem of “black box”, and the questions regarding the liability for harms caused by the systems relying on AI are common to all legal and policy considerations surrounding the themes of AI governance and regulation. They are relevant to European consumer law; but equally useful for those looking at AI from other normative angles. All three are topical issues, but pose different actual problems than usually assumed, stemming usually from conceptual confusions. In the next four sections, we look at issue specifically pertaining to the the European consumer law and policy.

## V. Asymmetry

Among public policy reasons for establishing a consumer protection framework the disparity of power between consumers and traders appears to be a central theme. Special consideration of consumer interests is generally explained by their weaker position with respect to the level of knowledge, resources, experience and the bargaining power.<sup>21</sup> This imbalance is seen as a source of potential concerns about efficiency, distributive justice and individual rights. In discussions among scholars and policymakers the junction between these three (and perhaps more) perspectives often remains debatable.<sup>22</sup> Ultimately, compound justifications appear to be most

---

<sup>21</sup> See, eg: cases C-137/08 *Pénzügyi Lizing*, ECLI:EU:C:2010:659, para 46 and C-105/17 *Kamenova*, ECLI:EU:C:2018:808, para 34.

<sup>22</sup> Cf S Grundmann, W Kerber and S Weatherill (eds) *Party Autonomy and the Role of Information in the Internal Market* (de Gruyter, 2001); F Rischkowsky and T Döring, ‘Consumer Policy in a Market Economy. Considerations from the Perspective of the Economics of Information, the New Institutional Economics as well as Behavioural Economics’ (2008) 31 *Journal of Consumer Policy* 285; I Benöhr, *EU Consumer Law and Human Rights* (Oxford University Press, 2013).

common and the emphasis placed on particular elements depends not only on the regulatory theory to which one subscribes, but also on the more general background (political, institutional) of the whole regulatory enterprise. This is well illustrated by the case of the European Union – a context which Hans Micklitz has explored in great detail. His analyses of the nexus between consumer law and the internal market policy, including his powerful conception of ‘access justice’,<sup>23</sup> have shed light on the specific dynamics in the European law and policy-making. At the same time, he has constantly remained on the lookout for the other, more socially oriented face of consumer policy at both national and European level.<sup>24</sup> As the ongoing socio-technological developments, including those related to AI, will continuously push us to reflect on the consumer protection instruments we have in place, the need to specify and agree on their underlying rationales will become even more pressing.

Much attention will likely remain, and rightly so, on the information asymmetry. It is a prominent aspect of the consumer’s weaker position and fits well both with the economic reasoning and with autonomy theory.<sup>25</sup> In simplified terms: without sufficient transparency of prices, quality and terms, consumers cannot make informed and therefore efficient choices, which in turn reduces suppliers’ incentives to compete on these parameters.<sup>26</sup> Consumers are perceived as rational actors capable of processing the relevant information and taking decisions that match their preferences and increase their utility. Such an approach was eagerly taken up by the EU legislator who has often prioritized the tools that serve ‘consumer empowerment’, such as disclosure duties.<sup>27</sup>

As already signalled before, however, the idea that consumers can improve the functioning of the market through their transactional decisions is not the only perspective that shapes the European law and policy-making. Legislation and case law on unfair terms in consumer contracts shows that consumers might be in need of more substantive protection when they are faced with contracts of adhesion.<sup>28</sup> The concept of ‘vulnerable consumers’ has also gradually emerged from the growing body of EU law, highlighting the fact that consumers who are not ‘reasonably well-informed and reasonably observant and circumspect’<sup>29</sup> do not necessarily have themselves to blame and may require a more targeted protection.<sup>30</sup> Later insights from behavioural economics

---

<sup>23</sup> H-W Micklitz, ‘Social Justice and Access Justice in Private Law’ (2011) EUI Working Paper LAW 2011/02. Available at: <http://cadmus.eui.eu/handle/1814/15706>.

<sup>24</sup> H-W Micklitz, ‘The Expulsion of the Concept of Protection from the Consumer Law and the Return of Social Elements in the Civil Law: A Bittersweet Polemic’ (2012) 35 *Journal of Consumer Policy* 283.

<sup>25</sup> Rischkowsky and Döring ‘Consumer Policy in a Market Economy’ (2008); O Ben-Shahar and CE Schneider, *More Than You Wanted To Know: The Failure of Mandated Disclosure* (Princeton University Press, 2014).

<sup>26</sup> I Ramsay, *Consumer Law and Policy: Text and Materials on Regulating Consumer Markets*, 3rd edn (Hart Publishing, 2012).

<sup>27</sup> H-W Micklitz, ‘The Consumer: Marketised, Fragmentised, Constitutionalised’ in D Leczykiewicz and S Weatherill (eds), *The Images of the Consumer in EU Law: Legislation, Free Movement and Competition Law* (Oxford, Hart Publishing, 2016).

<sup>28</sup> C Willett and M Morgan-Taylor, ‘Recognising the Limits of Transparency in EU Consumer Law’ in J Devenney and M Kenny (eds), *European Consumer Protection: Theory and Practice* (Cambridge University Press, 2012).

<sup>29</sup> See, eg: case C-210/96 *Gut Springenheide*, ECLI:EU:C:1998:369, para 31.

<sup>30</sup> N Reich, ‘Vulnerable Consumers in EU Law’ in Leczykiewicz and Weatherill (eds), *The Images of the Consumer* (2016).

have added further nuance to this picture, by pointing, among others, to certain systematic biases which affect consumer choice on a much wider scale.<sup>31</sup> The question which can now be asked is whether the changing technological landscape, and especially the developments in AI, are making this type of considerations more or less pronounced.

As should already be clear from sections II and IV, there will not be a single answer considering that AI is not a single ‘thing’. But first, it is helpful to take a step back and look at the process of digitalisation more generally and the consumers’ place in it. There are two contrasting viewpoints.<sup>32</sup> On the one hand, with an instantaneous access to the Internet and a variety of digital tools, contemporary consumers appear to be much better placed to find and compare information and exchange experience. Through this they can become more capable of assuming the role that neoclassical economists and, to a certain extent, the EU legislator have imagined for them.<sup>33</sup> On the other hand, due to the growing quantity of information, consumers often cannot meaningfully process and evaluate it on their own.<sup>34</sup> Both perspectives are true, even if they partially remain in tension. It is precisely this information paradox that has driven the development of a new generation of online intermediaries, such as search engines and social media, which allow consumers to navigate the information universe.

Although the solution came with a new unknown – as to how the filtering and personalisation actually take place and what they can lead to, this uncertainty has not weakened the demand for such tools and the pace of the technological progress. AI, or more precisely machine learning, is arguably the single most significant development in this process so far. It takes knowledge generation to the next level and makes it possible to create an entirely new value out of the information overload. The results seem ambiguous. Consumers gain access to potentially more relevant information about parts of the socio-economic landscape as well as to innovative products, such as personal digital assistants, which bring a new dimension to everyday convenience. At the same time, their continuous use of data-driven services and products (the line becomes blurred) contributes to the growing knowledge – and power – of their counterparts. What exactly are the underlying trade-offs often remains less clear, arguably as a result of both consumers’ cognitive limitations and different ‘black box’ problems (section III). This potentially adds to the already extensive insights traders have about the markets<sup>35</sup> and consumers – including

---

<sup>31</sup> B Schüller, ‘The Definition of Consumers in EU Consumer Law’ in Devenney and Kenny (eds), *European Consumer Protection* (2012); A-L Sibony and G Helleringer, ‘EU Consumer Protection and Behavioural Sciences: Revolution or Reform?’ in A Alemanno and A-L Sibony (eds), *Nudge and the Law: A European Perspective* (London, Hart Publishing, 2015).

<sup>32</sup> JK Winn, ‘Is Consumer Protection an Anachronism in the Information Economy?’ in JK Winn (ed) *Consumer Protection in the Age of the ‘Information Economy’* (Routledge, 2006).

<sup>33</sup> For a critical view see: C Coteanu, *Cyber Consumer Law and Unfair Trading Practices* (Routledge, 2005).

<sup>34</sup> More generally on the so-called information overload, see: G Howells, ‘The Potential and Limits of Consumer Empowerment by Information’ (2005) 32 *Journal of Law and Society* 349; Willett and Morgan-Taylor, ‘Recognising the Limits’ (2012); Ben-Shahar and Schneider, *More Than You Wanted to Know* (2014).

<sup>35</sup> M Weber, *Max Weber on Law in Economy and Society* (Harvard University Press, 1954).

their biases and misperceptions<sup>36</sup>. With AI, the possibilities of leveraging these factors to suppliers’ advantage can increase even more.

From this we can conclude that AI surely has a *potential* of increasing the asymmetry between consumers and traders. An important question is how far this change is actually taking place and whether it is part of an evolution, or rather a disruption. How broader is the knowledge which can now be inferred about consumers by looking at behaviours of countless other individuals? How does it translate into the existing regulatory rationales, eg does it affect consumers’ ability to ‘act rationally’? Recognizing that knowledge *about* consumers is now a crucial element of the traders’ informational advantage rightly brings the associations with data protection law to mind. It would be a mistake, however, to focus one’s attention solely on this domain. Rather, a combined reflection about data protection, consumer law, further neighbouring domains such as anti-discrimination law, and their underlying paradigms should be encouraged.<sup>37</sup>

## VI. AI Applications, Public Regulation and Private Commitment

There seems to be a broad consensus that machine learning as a science, or the enthusiasm which businesses express towards it, are in themselves neither good, nor evil and what needs to be assessed is a *specific* socio-technological practice (or AI application). Considering the asymmetrical dynamics mentioned before, as well as the transformative power associated with AI more generally, it is quite understandable, however, that remarkable advances in machine learning are causing mixed reactions.<sup>38</sup> It might be for these reasons that some of the leading companies engaged in the AI game have tried to position themselves as more than just profit-driven corporations. Both Facebook and Google pride themselves on their researchers’ publications in leading scientific journals along with their generous support of academic and social projects.<sup>39</sup> What is more, the industry leaders are increasingly interested in more than direct consumer interactions. One can observe a growing market not just for data, but also for machine learning

---

<sup>36</sup> JD Hanson and DA Kysar, ‘Taking Behavioralism Seriously: The Problem of Market Manipulation’ (1999) 74 *New York University Law Review* 630.

<sup>37</sup> W Kerber, ‘Digital Markets, Data, and Privacy: Competition Law, Consumer Law and Data Protection’ (2016) 11 *Journal of Intellectual Property Law & Practice* 856; N Helberger, FZ Borgesius and A Reyna ‘The Perfect Match? A Closer Look at the Relationship between EU Consumer Law and Data Protection Law’ (2017) 54 *Common Market Law Review* 1427; P Hacker, ‘Teaching Fairness to Artificial Intelligence: Existing and Novel Strategies against Algorithmic Discrimination under EU Law’ (2018) 55 *Common Market Law Review* 1143.

<sup>38</sup> See, eg: BEUC, ‘Automated Decision Making and Artificial Intelligence – A Consumer Perspective’ (2018) BEUC Position Paper. Available at: [http://www.beuc.eu/publications/beuc-x-2018-058\\_automated\\_decision\\_making\\_and\\_artificial\\_intelligence.pdf](http://www.beuc.eu/publications/beuc-x-2018-058_automated_decision_making_and_artificial_intelligence.pdf).

<sup>39</sup> See, eg: <https://research.fb.com/category/facebook-ai-research/>; <https://www.blog.google/outreach-initiatives/google-org/ai-social-good/>.

tools to be used by programmers.<sup>40</sup> This adds a further dimension to the discussions about the appropriate governance model for AI development and deployment.<sup>41</sup>

At this point legal scholars would welcome a categorization of some kind. What types of AI applications can we possibly expect to see in the market practice? When thinking about consumer markets, one can be tempted to distinguish between AI-based services designed for the use by consumers and those applied by businesses in marketing and customer relations. This is helpful to some extent. The former group would cover, for example, automated translation tools, recommender systems and personal assistants, while personalisation of commercial practices and contractual conditions or chatbots used in customer service would fall into the latter.<sup>42</sup> One needs to remember, however, that in reality the lines are more difficult to draw. This falls in line with the broader developments in the digital markets – monetization strategies can be diversified and indirect, different commercial relationships might be interconnected and the distinction between products and services often becomes blurred (with much of the exchange taking place on a long-term ‘as is’ basis). A given AI-based solution thus does not become more or less problematic due to the category in which it fits, but rather as a result of its overall context.

There might nevertheless be some guidelines of a more general nature that we would like AI developers to already consider at the programming stage. This view seems to be substantiated by the recent developments in the industry itself. ISO and IEEE have been active in developing standards for different aspects and use cases of AI applications.<sup>43</sup> As far as individual companies are concerned, most illustrative are probably the ‘AI principles’, which Google recently committed itself to follow in the hope to ‘lead by example’.<sup>44</sup> The motives for this move have already been questioned – and perhaps rightly so.<sup>45</sup> Equally intriguing are the narratives which accompany such bottom-up efforts – ones that typically pertain to ethical, not legal considerations. In fact, however, the connection between many of these initiatives and consumer law or the neighbouring domains, especially data protection and anti-discrimination law (Hans Micklitz would not approve of an overly sharp distinction), is readily apparent. In the field so complex and dynamic, getting right the relationship between public and private, and between the law, economy and ethics becomes particularly important.

What, then, are the main recurring themes in the ethical and legal discussions about AI? Certainly, a lot attention is devoted to discriminatory outcomes and their distributional effects.<sup>46</sup>

---

<sup>40</sup> See, eg: <https://facebook.ai/developers/tools>; <https://cloud.google.com/automl/>; <https://aws.amazon.com/machine-learning/>.

<sup>41</sup> Consider the idea of the European AI-on-demand platform as a possible response, see: European Commission, ‘Horizon 2020 Work Programme 2018-2020 Part 5.i Information and Communication Technologies’, 55. Available at: [http://ec.europa.eu/research/participants/data/ref/h2020/wp/2018-2020/main/h2020-wp1820-leit-ict\\_en.pdf](http://ec.europa.eu/research/participants/data/ref/h2020/wp/2018-2020/main/h2020-wp1820-leit-ict_en.pdf)

<sup>42</sup> Jabłowska, Kuziemski, Nowak, Micklitz, Pałka and Sartor, ‘Consumer Law and Artificial Intelligence’ (2018).

<sup>43</sup> See: <https://www.iso.org/committee/6794475.html>;  
[https://standards.ieee.org/news/2017/ieee\\_global\\_initiative.html](https://standards.ieee.org/news/2017/ieee_global_initiative.html)

<sup>44</sup> See: <https://www.blog.google/technology/ai/ai-principles/>

<sup>45</sup> In particular, the move is being linked to a controversial military contract concluded by the company with the US Department of Defense.

<sup>46</sup> T Zarsky, ‘Understanding Discrimination in the Scored Society’ (2014) 89 *Washington Law Review* 137; S Barocas and AD Selbst, ‘Big Data’s Disparate Impact’ (2016) 104 *California Law Review* 671; PT Kim, ‘Auditing

Considering that AI applications rely, at least partially, on human inputs (sourcing data, training algorithms), they are also susceptible to human imperfections. In contrast to single human errors and biases, however, once deployed, AI applications can carry large-scale implications and their shortcomings are not always easy to detect.<sup>47</sup> Since being prejudiced against underprivileged and minorities is one of the last thing major AI stakeholders would like to be associated with, preventing unfair bias seems to be an area of joint efforts. In its proudly unveiled principles, Google undertakes to avoid ‘unjust impacts on people, particularly those related to sensitive characteristics such as race, ethnicity, gender, nationality, income, sexual orientation, ability, and political or religious belief’. To be sure, this is not something we should be taking for granted or consider the problem solved. After all, until not long ago Facebook did not see anything wrong in allowing advertisers to target ads by race and ethnicity,<sup>48</sup> and the measures Google takes or intends to take to fulfil its declaration remain largely unclear.

Still, the growing number of the studies exploring technological measures for addressing the problem of unfair bias in machine learning gives grounds to hope that similar commitments can be more than just empty words.<sup>49</sup> Such a broad, multi-stakeholder interests in this area is certainly a welcome development. At the same time, as has been pointed out in the literature, hitherto attempts at formalizing fairness seem to rely on a limited perception of equality, understood as non-discrimination on the basis of a particular sensitive characteristics, such as gender or ethnicity.<sup>50</sup> This may be due to the current legal landscape, whose focus similarly remains on historically disadvantaged groups. From that perspective, certain characteristics mentioned by Google (eg income) appear to even go beyond the catalogues enshrined in law. An important question is, once again, when and how such legal-ethical conditions are being specified and implemented and who decides if they have been fulfilled. Finally, it is necessary to reflect on the more general limits of the sensitive characteristics (or protected classes) approach, as laid down in applicable legal frameworks. Arguably, AI systems can significantly transform the underlying populations over time as well as generate new societal segments, composed of multiple parameters, on a continuous basis.<sup>51</sup> This adds a new dimension to the existing line of research which investigates the problems of stigmatization and exclusion in a broader societal context.<sup>52</sup> We

---

Algorithms for Discrimination’ (2017) 166 *University of Pennsylvania Law Review Online* 189; Hacker, ‘Teaching Fairness’ (2018).

<sup>47</sup> C O’Neill, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (Allen Lane, 2016).

<sup>48</sup> See: <https://www.propublica.org/article/facebook-promises-to-bar-advertisers-from-targeting-ads-by-race-or-ethnicity-again>

<sup>49</sup> See: P Gajane and M Pechenizkiy, ‘On Formalizing Fairness in Prediction with Machine Learning’ (2017). Available at: <https://arxiv.org/abs/1710.03184>.

<sup>50</sup> R Binns, ‘Fairness in Machine Learning: Lessons from Political Philosophy’ (2018) 81 *Proceedings of Machine Learning Research* 149.

<sup>51</sup> LT Liu, S Dean, E Rolf, M Simchowitz and M Hardt, ‘Delayed Impact of Fair Machine Learning’ (2018) 80 *Proceedings of the 35th International Conference on Machine Learning* 3150; A Vedder and L Naudts, ‘Accountability for the Use of Algorithms in a Big Data Environment’ (2017) 31 *International Review of Law, Computers & Technology* 206.

<sup>52</sup> Eg intersectionality, see: L McCall, ‘The Complexity of Intersectionality’ (2005) 30 *Signs: Journal of Women in Culture and Society* 1771.

cannot also forget that legislative and judicial responses to the instances of discrimination tend to divide public opinion. From a governance perspective this implies that the questions raised by AI cannot often be ‘dealt with’ on a one-off basis. Addressing its impacts will rather be a long-term process, requiring engagement of multiple stakeholders and the civil society.

## VII. The Autonomy of Humans and the Autonomy of Machines

We have already expressed our scepticism about the gravity of the legal problems caused by machines acquiring ‘autonomous’ qualities. This does not mean that autonomy is entirely irrelevant in the context of AI. In fact, the contrary is true. However, as it is argued here, the focus should remain on the autonomy of humans.

This, at least theoretically, should bring us, lawyers, on a familiar ground. When thinking about autonomy in the legal context, ‘private autonomy’ – a traditional concept of private law and the main embodiment of the freedom of contract – is likely to be the first association to come in mind. Much of the scholarship over the last decades has been concerned with the conceptualization of both terms and their evolution.<sup>53</sup> The classical theory underscored the freedom of formally equal contract partners in shaping their relationships; as long as they *appear* to have agreed on a given transaction, the state should not intervene. Conceivable exceptions on the grounds of fairness have been those related to procedure, not substance and revolved around the so-called ‘vitiating factors’.<sup>54</sup> With the industrialisation and consumer society the discussion began to focus on the need of *overriding* private autonomy, in order to protect the party which is factually and structurally weaker.<sup>55</sup> However, as Hans Micklitz has shown in his work, there are also other ways of looking at autonomy<sup>56</sup>. This becomes particularly apparent when the EU law is considered – one in which these traditional theoretical underpinnings do not play a pivotal role and can be approached without prejudice. The concept of the ‘regulated autonomy’ tries to capture this multidimensional nature of autonomy in the European legal context.<sup>57</sup> Through this lens, European private law can be viewed as at the same time ‘establishing market freedoms, therefore increasing

---

<sup>53</sup> PS Atiyah, *The Rise and Fall of Freedom of Contract* (Oxford, Clarendon Press, 1979); H-W Micklitz, ‘On the Intellectual History of Freedom of Contract and Regulation’ (2015) 4 *Penn State Journal of Law and International Affairs* 1.

<sup>54</sup> P Cartwright, *Consumer Protection and the Criminal Law: Law, Theory, and Policy in the UK* (Cambridge University Press, 2001); E Mik, ‘The Erosion of Autonomy in Online Consumer Transactions’ (2016) 8 *Law, Innovation and Technology* 1.

<sup>55</sup> S Weatherill, ‘Justifying Limits to Party Autonomy in the Internal Market – EC Legislation in the Field of Consumer Protection’ in Grundmann, Kerber and Weatherill (eds), *Party Autonomy* (2001).

<sup>56</sup> For an equally relevant take on the regulatory dimension of autonomy in private law, see: S Grundmann, ‘On the Unity of Private Law. From a Formal to a Substance-Based Concept of Private Law’ (2010) 18 *European Review of Private Law* 1055.

<sup>57</sup> G Comparato and H-W Micklitz, ‘Regulated Autonomy Between Market Freedoms and Fundamental Rights in the Case Law of the CJEU’ in U Bernitz, X Groussot and F Schulyok (eds), *General Principles of EU Law and European Private Law* (Wolters Kluwer, 2013).

private autonomy’ and ‘set[ting] boundaries to this newly created autonomy’.<sup>58</sup> In doing so, the EU creates its own model of the freedom of contract characterised by the coexistence of these enabling and restricting (framing) dimensions.<sup>59</sup>

The challenges posed by AI are likely to make the ‘regulated autonomy’ an even more important perspective. Indeed, already today concerns are expressed in the literature that AI can lead to a much more powerful influence on consumer behaviours.<sup>60</sup> Crucially, the knowledge businesses may acquire through AI must be viewed in a broader context: with increasing digitalisation on a more general level it is currently possible to adjust not only the form and content of commercial messages to particular consumers, but also place and time.<sup>61</sup> Some of the recent studies began to explore the possibility of conceptualizing such personalised practices under Directive 2005/29/EC on unfair business-to-consumer commercial practices in the internal market (hereinafter UCPD).<sup>62</sup> With further advances in facial recognition, the importance of these questions will increase even more. A somewhat different line of research inquires about the impact of the delegated decision-making on autonomy and dignity.<sup>63</sup> This may seem like a different debate; yet, because of some of the elements of market transformation mentioned above, the line may be thinner than it looks at first sight. Just think about personal digital assistants and the vanishing divide between information and advertising.<sup>64</sup>

Against this background it seems pertinent to reflect on the parameters which one should take into account when deciding about the balance between the law’s enabling and restricting dimensions. As mentioned before, the EU seems to be primarily concerned with market integration. It recognizes the need for increased ‘consumer trust’ in the market in general, and in the internal market in particular, and fosters this trust through a variety of tools.<sup>65</sup> This, to some extent, remains in line with the regulatory rationale embedded in law and economics. To what

---

<sup>58</sup> H-W Micklitz and D Patterson, ‘From the Nation State to the Market: The Evolution of EU Private Law’ (2012) EUI Working Paper LAW 2012/15. Available at: <http://cadmus.eui.eu/handle/1814/22415>.

<sup>59</sup> Ibid.

<sup>60</sup> BEUC, ‘Automated Decision Making’ (2018).

<sup>61</sup> R Calo, ‘Digital Market Manipulation’ (2013) 82 *The George Washington Law Review* 995. Mik, ‘The Erosion of Autonomy’ (2016).

<sup>62</sup> 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. See: N Helberger, ‘Profiling and Targeting Consumers in the Internet of Things – A New Challenge for Consumer Law’ in R Schulze and D Staudenmayer (eds), *Digital Revolution: Challenges for Contract Law in Practice* (Nomos/Hart Publishing, 2015); Mik (n 54); M Sax, N Helberger and N Bol, ‘Health as a Means Towards Profitable Ends: mHealth Apps, User Autonomy, and Unfair Commercial Practices’ (2018) 41 *Journal of Consumer Policy* 1.

<sup>63</sup> R Brownsword, ‘From Erewhon to AlphaGo: For the Sake of Human Dignity, Should we Destroy the Machines?’ (2017) 9 *Law, Innovation and Technology* 117.

<sup>64</sup> ME Stucke and A Ezrachi, ‘How Digital Assistants Can Harm Our Economy, Privacy, and Democracy’ (2017) 32 *Berkeley Technology Law Journal* 1239; L Adam and H-W Micklitz, ‘Information, Beratung und Vermittlung in der digitalen Welt’ (2016) SVRV Working Paper Nr. 6. Available at: <http://www.svr-verbraucherfragen.de/dokumente/information-beratung-und-vermittlung-in-der-digitalen-welt/>

<sup>65</sup> Weatherill, ‘Justifying Limits’ (2001).

extent such a paradigm is indeed justified and to what extent it is fulfilled is a more complex matter, which we have touched upon in section V.

It is also important to stress the multidimensional character of the enabling-restricting dynamics, as this allows us escape false dichotomies. It would be too easy to say that with each consumer policy intervention, the autonomy of traders is (only) limited and the autonomy of consumers is (only) increased. What is more, different types of legal norms can be employed – for example information rules and rules affecting the content of legal relationships – all of them having different effects on the relevant actors. It thus seems that we need a deeper reflection on the nature of autonomy, the conditions under which it might be desirable for the law to step in, and the instruments with which this can be achieved.

Let us have a brief look at the instruments. Earlier reference to the UCPD may bring the traditional instruments employed against unfair commercial practices to mind. Depending on the Member State, these can be mainly actions for an injunction brought by consumer organisations or administrative decisions, including hefty sanctions, taken by a consumer protection agency.<sup>66</sup> So far Member States’ procedural autonomy in this field has remained rather broad, and the Court of Justice has not triggered the process of proceduralization comparable to the one known from its case law on Directive 93/13/EEC on unfair terms in consumer contracts.<sup>67</sup> Notably, the recent ‘New Deal for Consumers’ package, adopted by the European Commission, puts emphasis on a more harmonized enforcement.<sup>68</sup> The proposed amendments to the UCPD refer to ‘contractual and non-contractual remedies’, which should be available to consumers harmed by unfair commercial practices, as well as penalties, which should be added to the menu of sanctions available to the competent public authorities (administrative or judicial).<sup>69</sup> One can discuss at length whether the proposed approach is right or wrong, needed or not. The bottomline remains that all of these measures are inherently reactive. Assuming that commercial practices relying on AI could be qualified as infringements of the UCPD, such infringements would first need to be identified, which considering the personalised and elusive nature of commercial practices in the digital age may not be so easy to do, at least not with the current tools. For consumers, the ability to unilaterally terminate the contract could be a way of reversing the consequences of decisions taken while being subject to an undue influence. Similarly, however, if this remedy is work, it cannot be difficult to seek. Perhaps some inspiration could be found in the interface between public and private enforcement of competition law, although yet again, this is a complex matter, which cannot be addressed here in due detail.

There is, however, more to this debate than the right mix of remedies under UCPD. We have already emphasised the need for an integrated reflection into different domains. The previous

---

<sup>66</sup> H-W Micklitz and G Saumier (eds), *Enforcement and Effectiveness of Consumer Law* (Springer, 2018).

<sup>67</sup> Council Directive 93/13/EEC of 5 April 1993 on unfair terms in consumer contracts [1993] OJ L95/29; see, eg: cases C-388/13 *UPC Magyarország*, ECLI:EU:C:2015:225, paras 57-59 and C-109/17 *Bankia*, ECLI:EU:C:2018:735. See also: M Namysłowska and A Jabłonowska, ‘Member States Interests and the EU Law on Unfair B2C and B2B Practices’ in M Varju (ed), *Between Compliance and Particularism. Member State Interests and European Union Law* (Springer, 2019).

<sup>68</sup> COM(2018) 183 final.

<sup>69</sup> COM(2018) 185 final, Article 1.

section mentioned some aspects of anti-discrimination law, suggesting this is to be a field where the cooperation between public and private stakeholders is especially promising. With respect to the topic addressed here – the impact on consumers’ transactional decisions – the convergence of public and private interest is less pronounced, which possibly draws us towards more traditional tools of consumer protection, such as the ones discussed above. However, as we have seen, these are not without the problems of their own. More importantly perhaps, they seem to address the crux of the problem, the impact on consumers’ autonomy, only in a very indirect way – by providing for a fairly broadly phrased prohibition and ex post sanctioning in the hope to achieve a dissuasive effect. At the same time, in discussions about consumers and digitalisation, more and more attention appears to be directed to the ex ante perspective – what businesses can (or can be required to) do, also in a technological sense, to ensure that consumer autonomy is preserved. To be sure, this is not a dimension consumer law has not been preoccupied with so far. Quite the contrary, various pre-contractual information requirements bear witness to its importance.<sup>70</sup> However, in the more recent years, other tools, especially those of data protection law, are clearly gaining prominence. Specific requirements associated with consent, transparency and control over the use of data, or privacy by design and by default are some of the most emphasised themes.<sup>71</sup> With all public attention devoted to the General Data Protection Regulation (hereinafter: GDPR)<sup>72</sup> they have been extensively explored and publicized. The impact on the market practice is also visible. Now is time to look at all these tools and their practical operation in still more detail and connect the respective debates.

## VIII. Instead of Conclusions: Empowerment

All the discussion up to this point has been concerned with AI-powered and big data fueled tools developed and used by big business and/or state exercising its policing power in criminal justice or tax collection. This is because, as of late 2018, actual usage of artificial intelligence is almost exclusively the domain of the corporations and the state. The reason is simple: creation of new AI applications is resource-intensive, requires substantial funding, and relies on the existence/creation of enormous data sets, readily available to online giants and law enforcement agencies. However, it does not have to be that way.

---

<sup>70</sup> See, eg the discussions surrounding the adoption of 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 [2011] OJ L304/64.

<sup>71</sup> See especially guidelines and opinions of Article 29 Working Party and the European Data Protection Supervisor.

<sup>72</sup> 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.

In the spring of 2016, one of us was eating an ice cream and enjoying the sun in the gardens of Villa Schifanoia in Florence, when Hans Micklitz – passing by – stopped and said: “You know what? I think we should build a software that will read online terms of service and highlight unfair contractual clauses. These contracts are way too long for anyone to read; and consumer law is a mess. Everyone would be way better off if computers could do that”. The initial response was a list of reasons why this cannot work: we are not programmers, natural language is complex, law cannot be formalized... but Hans just raised his hand, stopped the outpour of scepticism, and said: “instead of telling me why this cannot be done, make a list of conditions for making it possible. Let’s try”. Half a year later, a prototype of that software was finished,<sup>73</sup> and a year later an article with the theory of consumer empowerment using technology was published.<sup>74</sup> That piece of software did not rely on artificial intelligence yet; but the project did not stop there. In the summer of 2017 a new project – CLAUDETTE – got launched, with the ambition of automating the analysis of online terms of service<sup>75</sup> and privacy policies<sup>76</sup> using machine learning. All this within the paradigm of empowering the individuals and the civil society using technology. This is Hans Micklitz in action. He is both a thinker, and a doer.

Within this spirit, we want to conclude the Chapter with the message: artificial intelligence can and should be used to empower the individuals and the civil society. Clearly, the challenges it raises are real, and should be tackled by law and/or policy. However, legal and political response to the developments of artificial intelligence should not limit itself to addressing these challenges only. Law should also, if not even more, concentrate on enabling creations of AI-powered systems to help empower the people. What would that mean?

In the domain of consumer law, individuals and civil society organisations enjoy numerous rights and competences. Countless information obligations give consumer a right to access information about products and services they use. Unfair contractual terms rules give consumers and their organizations standing to challenge certain clauses in boilerplate contracts that create an imbalance of rights and duties for the detriment to consumers. GDPR contains numerous information duties for the data controllers, and so consumers (data subject) have a right to know who uses what data about them and for what purpose. Unfair commercial practices legislation enables us to challenge misleading and/or aggressive ads and offers. We have a lot of rights. But who has time and power to make use of them? Put provocatively, the problem of consumers and NGOs is not the lack of the legal powers, but rather too little resources and factual abilities to make use of them.

---

<sup>73</sup> The Reader can download it at <http://utermsoftware/download/>.

<sup>74</sup> H-W Micklitz, P Palka and Y Panagis, ‘The Empire Strikes Back: Digital Control of Unfair Terms of Online Services’ (2017) 40 *Journal of Consumer Policy* 367.

<sup>75</sup> M Lippi, P Palka, G Contissa, F Lagioia, H-W Micklitz, G Sartor, P Torroni, ‘CLAUDETTE: an Automated Detector of Potentially Unfair Clauses in Online Terms of Service’ (2018). Available at: <https://arxiv.org/abs/1805.01217>

<sup>76</sup> G Contissa, K Docter, F Lagioia, M Lippi, H-W Micklitz, P Palka, G Sartor, P Torroni, ‘Claudette Meets GDPR: Automating the Evaluation of Privacy Policies Using Artificial Intelligence’ (2018). Available at: [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3208596](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3208596).

That is where AI can be of help. Reading of contracts and privacy policies can be automated. Analysis of what apps actually do can be automated. Communication with businesses and consumer protection bodies can be automated. We could imagine a world in which consumers are equipped with apps that read the agreements they accept for them, detect unfair clauses, and notify the companies and the enforcers. We could live in a world where ads are assessed for unfairness. We could live in a world where exercise of many rights – like to object to data processing – is automatic. Big part of the work initiated by Hans Micklitz within this strand concentrated on outlining the theoretical preconditions for these developments, and presenting proof for of concept. On the technical side, we are not far from that world. The problem again is the resources. What is needed as of today is an interdisciplinary collaboration of engineers and lawyers, and financial means for engaging in costly development processes.

To paraphrase president Kennedy, who – as Hans Micklitz often reminds us – was one of the major figures in the development of the consumer law: ask not only what consumer law can do with artificial intelligence, ask what artificial intelligence can do for consumer law!

## Bibliography

Adam, L and Micklitz, H-W, ‘Information, Beratung und Vermittlung in der digitalen Welt’ (2016) SVRV Working Paper Nr. 6.

Alpaydin, E, *Machine Learning. The New AI* (MIT Press, 2016).

Atiyah, PS, *The Rise and Fall of Freedom of Contract* (Oxford, Clarendon Press, 1979).

Barocas, S and Selbst, AD, ‘Big Data’s Disparate Impact’ (2016) 104 *California Law Review* 671.

Benöhr, I, *EU Consumer Law and Human Rights* (Oxford University Press, 2013).

Ben-Shahar, O and Schneider, CE, *More Than You Wanted To Know: The Failure of Mandated Disclosure* (Princeton University Press, 2014).

BEUC, ‘Automated Decision Making and Artificial Intelligence – A Consumer Perspective’ (2018) BEUC Position Paper.

Binns, R, ‘Fairness in Machine Learning: Lessons from Political Philosophy’ (2018) 81 *Proceedings of Machine Learning Research* 149.

Brownsword, R, ‘From Erewhon to AlphaGo: For the Sake of Human Dignity, Should we Destroy the Machines?’ (2017) 9 *Law, Innovation and Technology* 117.

Calo, R, ‘Digital Market Manipulation’ (2013) 82 *The George Washington Law Review* 995.

Cartwright, P, *Consumer Protection and the Criminal Law: Law, Theory, and Policy in the UK* (Cambridge University Press, 2001).

Contissa, G, Lagioia, F and Sartor, G, ‘The Ethical Knob: Ethically-Customisable Automated Vehicles and the Law’ (2017) 25 *Artificial Intelligence and Law* 365.

Contissa, G, Docter, K, Lagioia, F, Lippi, M, Micklitz, H-W, Pałka, P, Sartor, G, Torroni, P, ‘Claudette Meets GDPR: Automating the Evaluation of Privacy Policies Using Artificial Intelligence’ (2018). Available at: [www.ssrn.com](http://www.ssrn.com).

Comparato, G and Micklitz, H-W, ‘Regulated Autonomy Between Market Freedoms and Fundamental Rights in the Case Law of the CJEU’ in U Bernitz, X Groussot and F Schulyok (eds), *General Principles of EU Law and European Private Law* (Wolters Kluwer, 2013).

Coteanu, C, *Cyber consumer law and unfair trading practices* (Routledge, 2005).

Cristianini, N, ‘The Road to Artificial Intelligence: A Case of Data Over Theory’ (2016) *The New Scientist*.

Gajane, P and Pechenizkiy, M, ‘On Formalizing Fairness in Prediction with Machine Learning’ (2017). Available at: [www.arxiv.org](http://www.arxiv.org).

Grundmann, S, ‘On the Unity of Private Law. From a Formal to a Substance-Based Concept of Private Law’ (2010) 18 *European Review of Private Law* 1055.

Grundmann, S, Kerber, W and Weatherill, S (eds), *Party Autonomy and the Role of Information in the Internal Market* (de Gruyter, 2001).

Hacker, P, ‘Teaching Fairness to Artificial Intelligence: Existing and Novel Strategies against Algorithmic Discrimination under EU Law’ (2018) 55 *Common Market Law Review* 1143.

Hanson, JD and Kysar, DA, ‘Taking Behavioralism Seriously: The Problem of Market Manipulation’ (1999) 74 *New York University Law Review* 630.

Helberger, N, ‘Profiling and Targeting Consumers in the Internet of Things – A New Challenge for Consumer Law’ in R Schulze and D Staudenmayer (eds), *Digital Revolution: Challenges for Contract Law in Practice* (Nomos/Hart Publishing, 2015).

Helberger, N, Borgesius, FZ and Reyna, A, ‘The Perfect Match? A Closer Look at the Relationship between EU Consumer Law and Data Protection Law’ (2017) 54 *Common Market Law Review* 1427

Howells, G, ‘The Potential and Limits of Consumer Empowerment by Information’ (2005) 32 *Journal of Law and Society* 349;

Jabłonowska, A, Kuziemski, M, Nowak, AM, Micklitz, H-W, Pałka, P, and Sartor, G, ‘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’ (2018) EUI Working Paper LAW 2018/11.

Kerber, W, ‘Digital Markets, Data, and Privacy: Competition Law, Consumer Law and Data Protection’ (2016) 11 *Journal of Intellectual Property Law & Practice* 856

Kim, PT, ‘Auditing Algorithms for Discrimination’ (2017) 166 *University of Pennsylvania Law Review Online* 189.

Kurzweil, R, *The Age of Intelligent Machines* (Cambridge, MA, MIT Press, 1990).

Lippi M, Pałka P, Contissa G, Lagioia F, Micklitz H-W, Sartor G, Torroni P, ‘CLAUDETTE: an Automated Detector of Potentially Unfair Clauses in Online Terms of Service’ (2018). Available at: [www.arxiv.org](http://www.arxiv.org).

Liu, LT, Dean, S, Rolf, E, Simchowit, M and Hardt, M, ‘Delayed Impact of Fair Machine Learning’ (2018) 80 *Proceedings of the 35th International Conference on Machine Learning* 3150.

McCall, L, ‘The Complexity of Intersectionality’ (2005) 30 *Signs: Journal of Women in Culture and Society* 1771.

Micklitz, H-W, ‘Social Justice and Access Justice in Private Law’ (2011) EUI Working Paper LAW 2011/02.

Micklitz, H-W, ‘The Expulsion of the Concept of Protection from the Consumer Law and the Return of Social Elements in the Civil Law: A Bittersweet Polemic’ (2012) 35 *Journal of Consumer Policy* 283.

Micklitz, H-W, ‘On the Intellectual History of Freedom of Contract and Regulation’ (2015) 4 *Penn State Journal of Law and International Affairs* 1.

Micklitz, H-W, ‘The Consumer: Marketised, Fragmentised, Constitutionalised’ in D Leczykiewicz and S Weatherill (eds), *The Images of the Consumer in EU Law: Legislation, Free Movement and Competition Law* (Oxford, Hart Publishing, 2016).

Micklitz, H-W and Patterson, D, ‘From the Nation State to the Market: The Evolution of EU Private Law’ (2012) EUI Working Paper LAW 2012/15.

Micklitz, H-W and Saumier, G (eds), *Enforcement and Effectiveness of Consumer Law* (Springer, 2018).

Mik, E (2016) ‘The Erosion of Autonomy in Online Consumer Transactions’ (2016) 8 *Law, Innovation and Technology* 1.

Namysłowska, M and Jabłonowska, A, ‘Member State Interests and the EU Law on Unfair B2C and B2B Practices’ in M Varju (ed), *Between Compliance and Particularism. Member State Interests and European Union Law* (Springer, 2019).

O’Neill, C, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (Allen Lane, 2016).

Pałka, P, Jabłonowska, A, Micklitz, H-W and Sartor, G, ‘Before Machines Consume the Consumers: High-Level Takeaways from the ARTSY Project’ (2018) EUI Working Paper LAW 2018/12.

Ramsay, I, *Consumer Law and Policy: Text and Materials on Regulating Consumer Markets*, 3rd edn (Hart Publishing, 2012).

Reich, N, ‘Vulnerable Consumers in EU Law’ in D Leczykiewicz and S Weatherill (eds), *The Images of the Consumer in EU Law: Legislation, Free Movement and Competition Law* (Oxford, Hart Publishing, 2016).

Rischkowsky, F and Döring, T, ‘Consumer Policy in a Market Economy. Considerations from the Perspective of the Economics of Information, the New Institutional Economics as well as Behavioural Economics’ (2008) 31 *Journal of Consumer Policy* 285.

Russell, SJ and Norvig, P, *Artificial Intelligence: A Modern Approach*, 3rd edn (Pearson Education Limited, 2016).

Sax, M, Helberger, N and Bol, N, ‘Health as a Means Towards Profitable Ends: mHealth Apps, User Autonomy, and Unfair Commercial Practices’ (2018) 41 *Journal of Consumer Policy* 1.

Schüller, B, ‘The Definition of Consumers in EU Consumer Law’ in Devenney and Kenny (eds), *European Consumer Protection: Theory and Practice* (Cambridge University Press, 2012).

Sibony, A-L, and Helleringer, G, ‘EU Consumer Protection and Behavioural Sciences: Revolution or Reform?’ in A Alemanno and A-L Sibony (eds), *Nudge and the Law: A European Perspective* (London, Hart Publishing, 2015).

Silver, D, Huang, A, Maddison, CJ, Guez, A, Sifre, L, van den Driessche, G, Schrittwieser, J, Antonoglou, I, Panneershelvam, V, Lanctot, M, Dieleman, S, Grewe, D, Nham, J, Kalchbrenner, N, Sutskever, I, Lillicrap, T, Leach, M, Kavukcuoglu, K, Graepel, T and Hassabis, D, ‘Mastering the Game of Go with Deep Neural Networks and Tree Search’ (2016) 529 *Nature* 484.

Silver, D, Schrittwieser, J, Simonyan, K, Antonoglou, I, Huang, A, Guez, A, Hubert, T, Baker, L, Lai, M, Bolton, A, Chen, Y, Lillicrap, T, Hui, F, Sifre, L, van den Driessche, G, Graepel, T and Hassabis D, ‘Mastering the Game of Go without Human Knowledge’ (2017) 550 *Nature* 354.

Stucke, ME and Ezechia, A, ‘How Digital Assistants Can Harm Our Economy, Privacy, and Democracy’ (2017) 32 *Berkeley Technology Law Journal* 1239.

Vedder, A and Naudts, L, ‘Accountability for the Use of Algorithms in a Big Data Environment’ (2017) 31 *International Review of Law, Computers & Technology* 206.

Weatherill, S, ‘Justifying Limits to Party Autonomy in the Internal Market – EC Legislation in the Field of Consumer Protection’ in S Grundmann, W Kerber and S Weatherill (eds), *Party Autonomy and the Role of Information in the Internal Market* (de Gruyter, 2001).

Weber, M, *Max Weber on Law in Economy and Society* (Harvard University Press, 1954).

Willett, C and Morgan-Taylor, M, ‘Recognising the Limits of Transparency in EU Consumer Law’ in J Devenney and M Kenny (eds), *European Consumer Protection: Theory and Practice* (Cambridge University Press, 2012).

Winn, JK, ‘Is Consumer Protection an Anachronism in the Information Economy?’ in JK Winn (ed) *Consumer Protection in the Age of the ‘Information Economy’* (Routledge, 2006).

Winston, PH, *Artificial Intelligence*, 3rd edn (Addison-Wesley, 1992).

Zarsky, T, ‘Understanding Discrimination in the Scored Society’ (2014) 89 *Washington Law Review* 137.