

# Artificial Intelligence: Challenges for EU Citizens and Consumers

## KEY FINDINGS

This briefing addresses the regulation of artificial intelligence (AI), namely, how to ensure that AI benefits citizens and communities, according to European values and principles. Focusing on data and consumer protection, it presents risks and prospects of the applications of AI, it identifies the main legal regimes that are applicable, and examines a set of key legal issues.

Key observations on the development of AI and its impact on citizens and consumers include the following:

- The convergence between data-driven AI and the global data-processing infrastructure raises risks for citizens as the power of AI can be harnessed for surveillance and manipulation. The combined powers of AI and big data can restrict users' options, influence their opinions and manipulate them into making choices that do not serve their best interests.
- Both legal regulation and social empowerment are needed to ensure that AI is developed and deployed in ways that preserve and enhance individual interests and the social good.
- Legal regulation has to focus on first principles, including individual rights and social goals, as well as on existing regulatory frameworks, such as data protection, consumer protection and competition law.
- Multiple legally relevant interests are affected by AI, such as data protection (lawful and proportionate processing of personal data, subject to oversight), fair algorithmic treatment (not being subject to unjustified prejudice resulting from automated processing), transparency and explicability (knowing how and why a certain algorithmic response has been given or a decision made), protection from undue influence (not being misled, manipulated, or deceived).

Key issues in the consumer domain to be addressed include the following:

- The extent to which algorithmic price discrimination is acceptable in online markets should be clarified.
- Inacceptable practices in targeted advertising and nudging directed to consumers should be defined and addressed.
- Discrimination in ads delivery should be countered.
- Citizens and consumers should be provided with effective ways to turn off personalisation.
- The development and deployment of AI tools that empower citizens and consumers and civil society organisations should be incentivised.



## The concept of AI

The broadest definition of artificial intelligence (AI) characterises it as the attempt to build machines that “perform functions that require intelligence when performed by people”<sup>1</sup>. Developing appropriate policies and regulations for AI is a priority for the European Union (EU), since AI has become a powerful driver of social transformation, reshaping individual lives and interactions as well as economical and political organisations. AI brings huge opportunities (for development, sustainability, health, knowledge, etc.) as well as significant risks (unemployment, discrimination, exclusion, etc.). Thus, on the one hand we must support research in AI, as well as the development and deployment of useful AI applications, and on the other hand we must provide an appropriate legal and ethical framework based on European values so as to ensure that AI contributes to the social good<sup>2</sup>. We should ensure that AI is used in a way that is consonant with human nature and fosters its potentialities, and that AI is neither underused, with opportunity costs, nor overused or misused, with the resulting risks<sup>3</sup>. In analysing the social impacts of AI, we must consider that AI is not a single technology, but rather a vast set of diverse approaches and technologies, which to different degrees and in different ways display intelligent behaviour in many different contexts. Some of these technologies are disruptive of existing social arrangements, while others fit smoothly into them.

AI has gone through a number of ups and downs since its beginnings in the 1950s, overly optimistic expectations being followed by bitter disillusion<sup>4</sup>. However, there is no doubt that in recent years AI has been hugely successful. A solid interdisciplinary background has been constructed for AI research: the original core of computing, engineering, mathematics, and logic has been extended with models and insights from a number of other disciplines, such as statistics, economics, linguistics, the neurosciences, philosophy and law. Furthermore, an array of successful applications has been built which have already entered people’s lives: voice, image, and face recognition; automated translation; document analysis; question-answering; games; high speed trading; industrial and home robotics; autonomous vehicles; etc.

Two kinds of AI are often distinguished. On the one hand, AI research may lead to “artificial general intelligence”, also called “strong AI”, namely, to the creation of computer systems that exhibit most of the human cognitive skills, at the human level and also at a superhuman level. On the other hand, AI research pursues the more modest objective of constructing “artificial specialised intelligence”, also called “narrow AI”, i.e., systems capable of satisfactorily carrying out single specific tasks requiring intelligence.

The possible future emergence of “artificial general intelligence” already raises serious concerns. Some researchers have argued for the need to anticipate the “existential threats” resulting from superintelligent AI systems by adopting measures meant to prevent the emergence of superintelligence or else to ensure that it will be directed towards human-friendly outcomes<sup>5</sup>. Though the risks related to the emergence of “strong AI” should not be underestimated (this is, on the contrary, a most salient problem to be addressed in the future), it seems too early now to address it at the policy level. Indeed, the creation of an “artificial general intelligence” lies decades ahead. A greater experience with advanced AI and a broader debate is needed before we can fully understand both the extent and the proximity of the risks of “strong AI”, and the best ways to address them. Thus, in this briefing, only the multiple realisations of “narrow AI” (“artificial specialised intelligence”) will be addressed, which already raise a number of legal and social issues.

## AI and the global data-infrastructure

The success of AI is linked to a change in the leading paradigm in AI research and development. Until a few decades ago, it was generally assumed that in order to develop an intelligent system, humans had to provide a formal representation of the relevant knowledge (usually expressed through a combination of rules and concept), coupled with algorithms serving to make inferences out of such knowledge. More recently, the focus has shifted to the possibility of applying machine-learning algorithms to vast masses of data: given a vast set of examples of correct or interesting behaviour, collected in a dataset, a system can learn by itself

how to address new cases on the basis of the examples, without the need for human modelling of the relevant knowledge<sup>6</sup>. For many applications, the learning set consists of information about individual or social behaviours, namely, of personal data. Thus, AI has become hungry for personal data, and this hunger has spurred data collection<sup>7</sup>.

In fact, the development of data-driven AI applications both presupposes and stimulates the availability of huge data-sets, the so-called big data<sup>8</sup>. A momentous process of digitisation has indeed preceded most applications of AI, resulting from the fact data flows are produced in all domains where computing is deployed<sup>9</sup>: from computer mediated economic transactions (as in e-commerce)<sup>10</sup>, from sensors monitoring and providing input to physical objects (e.g., cars, smart home devices), from the workflows of economic and governmental activities (e.g., banking, transportation, taxation), from surveillance devices (e.g., street cameras), and, last but not most significant, from non-market activity of individuals (e.g., Internet access, search, social networking). These data flows have been integrated into a global interconnected data-processing infrastructure, centred on, but not limited to, the Internet. The infrastructure provides a universal medium to communicate, to access data, and to deliver any kind of private and public services. It enables citizens to shop, use banking and other services, pay taxes, get benefits, access information and knowledge, and build social networks. Algorithms mediate citizens' access to content and services, selecting for them information and opportunities.

Thanks to the integration of AI into the global data-processing infrastructure, a lot of good can be provided: overcoming the information overload, world-wide generation and distribution of knowledge and solutions, economic efficiency, wealth creation, individualised private and public services, environmentally-friendly management of utilities and logistics, support for transparency, discovering and remedying biases and discriminations. AI enables researchers to discover unexpected correlations and develop evidence-based models, doctors to provide personalised healthcare, businesses to detect market trends and make more efficient decisions, consumers to make informed choices and obtain personalised services, public authorities to identify risks, prevent harm, better manage public goods (e.g., the environment) and coordinate citizens' actions (e.g., traffic).

However, this convergence of AI and data-processing also leads to serious risks for individuals and society. Profit-making actors can harness the power of AI (as applied to vast masses of data) to pursue legitimate economic goals in ways that are harmful to individuals and society: they can subject citizens to pervasive surveillance, unduly restrict their options and the information they access, and manipulate them into making choices that do not serve their best interests<sup>11</sup>. In fact, many Internet companies (such as the major platforms for user-generated contents) operate in two-sided markets: their main services (e.g., search engines, social networking) are offered to individual consumers, but revenue comes from advertisers and persuaders (e.g., in political campaigns). This means not only that all information that may be useful for targeted advertising will be collected and used for this purpose, but also that platforms will use any means to capture users, so that they can be exposed to ads and attempts to persuade them. As noted, users may be captured by giving them information they like, or which accords with their preferences, thereby exploiting their confirmation biases<sup>12</sup>, which may lead to polarisation and fragmentation in the public sphere, and to the proliferation of sensational and fake news.

Similarly, governments may use AI for legitimate political and administrative purposes (e.g., efficiency, cost reduction, improved services), but also to anticipate and control citizens' behaviour in ways that restrict individual liberties and interfere with democratic processes. It has indeed been argued that both a surveillance capitalism<sup>13</sup> and a surveillance state<sup>14</sup> are developing and converging. The PRISM program in the US and the social credit system being developed in China are worthy of note as significant examples of disproportionate, pervasive mass surveillance merging data collected by private companies and governmental agencies.

## Legal regulation and social empowerment

Valuable sociotechnical practices around the use of AI need to be promoted, whilst asking ourselves how we can ensure that the development and deployment of AI tools takes place in contexts (inclusive of technologies, human skills, organisational structures, and norms) where individual interests and the social good are both preserved and enhanced. Here this briefing will focus on two aspects: legal regulation and social empowerment.

With regards to legal regulation, we need to focus not only on existing regulatory frameworks, but also on first principles, given that the available rules may fail to provide appropriate solutions and direction to citizens and legal decision-makers. First principles include fundamental rights and social values at both the ethical and the legal level. A high-level synthesis of the ethical framework for AI is provided by the recent AI4People document, which describes as follows the risks and opportunities provided by AI<sup>15</sup>:

- enabling human self-realisation, without devaluing human abilities;
- enhancing human agency, without removing human responsibility; and
- cultivating social cohesion, without eroding human self-determination.

Moving from ethics to law, AI is able to both promote and endanger not only the fundamental rights to privacy and data protection but also other rights included in the European Charter, starting with dignity, and including freedom of thought, conscience, and religion, freedom of expression and information, freedom of assembly and of association, freedom to choose an occupation, equality, non-discrimination, and consumer protection. Many fundamental social goals (such as welfare, competition, efficiency, advancement in science, art and culture, cooperation, civility, and democracy) are also at stake in the deployment of AI.

As AI is affecting many aspects of citizens' individual and social lives, it falls under the scope of different sectorial legal regimes, especially data protection law, consumer protection law, and competition law. As has been observed by the European Data Protection Supervisor (EDPS) in its Opinion 8/18 on the legislative package "A New Deal for Consumers", there is synergy between the three regimes. Consumer and data protection law share the common goals of correcting imbalances of informational and market power, and, along with competition law, they contribute to ensuring that people are treated fairly. However, the EDPS detected a possible conflict between the idea in the General Data Protection Regulation (GDPR) that people have a fundamental right to data protection and the view, emerging from some provisions in the new consumer package, that personal data are a tradable property. Similarly, the need to limit the market power of the biggest players and to preserve consumer choice in the long run can come into in conflict with some approaches to competition law, such as those that focus on enhancing consumer welfare in the present.

## Interests/rights at stake and risks from a consumer/user perspective

Consumers are highly at risk in a context where AI technologies are deployed in the service of business entities and are used to influence consumer purchasing and other consumer behaviour. In this context, a number of legally relevant interests may be affected.

First, there is the interest in data protection, namely, the interest in a lawful and proportionate processing of personal data subject to oversight. This is hardly compatible with an online environment where every action is tracked, and the resulting data is used to extract further information about the individuals concerned, beyond their control, and to process this information in ways that may run counter to their interests.

The processing of personal data through AI systems may also affect citizens' interest in fair algorithmic treatment, namely, the interest in not being subject to unjustified prejudice resulting from automated processing. In particular, consumers may be presented with higher prices targeted specifically at them, or they may be kept uninformed or be prevented from finding out about opportunities that would be significant to them (e.g., loans, rental contracts, etc.).

The possibility of algorithmic unfairness, as well as the need to keep the processing of personal data under control and to understand (and possibly challenge) the reasons for determinations that affect individuals, give rise to a concern from an algorithmic transparency/explicability standpoint, namely, citizens want to know how and why a certain algorithmic response has been given or a decision made, so as "to understand and hold to account the decision-making processes of AI"<sup>16</sup>. This concerns significant decisions, by public or private powers (e.g., access to jobs and positions, granting loans, allocating benefits, imposing sanctions). However, transparency/explicability should also be available when, on the basis of profiling, citizens are the object of a set of micro-decisions that individually are not very important but which, on the whole, may have a substantial impact on them. With regards to both significant decisions and on profile-based micro-decisions, individual autonomy is affected when citizens interact with black boxes<sup>17</sup>, whose functioning is not accessible to them, and whose decisions remain unexplained and thus unchallengeable<sup>18</sup>.

As AI systems have access to a huge amount of information about individuals and about people similar to them, they can effortlessly use this information to elicit desired behaviour, for purposes such citizens may not share, possibly in violation of fiduciary expectations they have toward the organisation that is using the AI system in question<sup>19</sup>. Thus, individuals have an interest in not being misled or manipulated by AI systems, but they also have an interest in being able to trust such systems, knowing that they will not profit from the people's exposure (possibly resulting from personal data). Reasonable trust is needed for individuals not to waste their limited and costly cognitive capacities in trying to fend off AI systems' attempts to mislead and manipulate them. In particular, consumers may be manipulated when their decisions are guided by a digital environment that makes certain valuable choices less accessible to them or directs their limited cognitive capacities towards outcomes they might regret or would have regretted had they been provided with better knowledge. Consumers are in a weak position when facing automated persuaders, which have access to a huge amount of knowledge, can effortlessly deploy unlimited computational power, and can frame the choice and information environment that is available to consumers. Moreover, the need to capture users in a two-sided market may lead to the phenomena we call filter bubbles or echo chambers<sup>20</sup>, namely, to the fact that users are provided with information on the basis of the extent to which such information is likely to capture them and stimulate a desired behaviour in them. The very data collected for the purpose of targeted advertising (e.g., by tracking users when they make searches or click on webpages, or by harvesting the data provided by users themselves such as their likes or social connections) can be used to influence political opinions and choices, as happened in recent elections, as evidenced by the Cambridge Analytica case.

Citizens and consumers also have an indirect interest in fair algorithmic competition, i.e., in not being subject to market-power abuses resulting from exclusive control over masses of data and technologies. This is of direct concern to competitors, but the lack of competition may negatively affect consumers, too, by depriving them of valuable options and restricting their sphere of action.

In the following subsections, this briefing will provide a short review of some significant issues in the consumer domain which are related to the interests just mentioned, and which raise open legal and policy issues.

## Issue: Price Discrimination

AI supports suppliers not only in directly targeting ads and recommendations to consumers, but in presenting them with individualised prices, namely, in offering to each consumer an approximation of the highest price point that consumer may be able or willing to pay. Lawyers, ethicists, and economists have been debating the legitimacy of price discrimination in the consumer domain<sup>21</sup>. Certain markets, such as credit or insurance, do indeed operate on cost structures based on risk profiles correlated with features distinctive to individual consumers, suggesting that it may be reasonable to offer different prices (e.g., interest rates) to different consumers. Should we allow price discrimination in other cases, too, for instance on the basis of the ability of different consumers to pay or even by exploiting their needs and vulnerabilities?

Various normative standards could be made available in making this assessment. We may refer to consumer protection to determine whether price discrimination, in a certain context, may be deemed unfair or discriminatory or based on aggressive or misleading advertising. We may look to the GDPR and ask whether it may count as an automated decision under its Article 22 or, more generally, whether this kind of processing serves “legitimate interests pursued by the controller or by a third party” (interests not “overridden by the interests or fundamental rights and freedoms of the data subject” under Article 6). Finally, even competition law can be at stake, as opaque price discrimination may undercut competition among suppliers in a marketplace in which only some, but not others, can take advantage of access to the data on which price discrimination can be based.

## Issue: Targeted Advertising/Nudging

AI systems can deliver to each consumer ads that are most likely to result in the desired purchasing behaviour given the available data. The ads being served may answer the consumers’ interests and help consumers navigate through the overwhelming array and variety of online markets. However, they may also exploit consumer vulnerabilities (e.g., predatory loans offered to people in financial distress, gambling offers made to gambling addicts, psychoactive drugs marketed to depressed people). A system which exploits users’ vulnerabilities may not be “intentionally malicious”: it may just have learned to send a certain kind of message to a person (e.g., posts from a gambling site to an addicted gambler, addictive fashion ads to a compulsive shopper), merely by having been programmed to learn what behaviour is likely to generate higher sales. Nudges meant to induce desired behaviour patterns can be exercised in subtle ways, by framing the contexts in which individuals are going to exercise their choices and selecting what information is made available to each of them<sup>22</sup>.

Different principles can be put forward on this unlawful or borderline practice. According to data protection law, we may ask whether exploiting people’s vulnerabilities to induce them to make choices they will regret can be viewed, on balance, as a legitimate use of personal data. According to consumer protection law, we may similarly consider whether this should count as aggressive advertising. A more general issue is raised by the possibility that AI systems should operate in ways that, while benefitting their deployers, cause individual and social harm. What are the duties of care (or obligations of solidarity) that apply to those that deploy AI to enhance their own capacity for action? Deploying AI in ways that may give some an advantage over a counterpart may be permissible when the competition among the interests at stake yields socially beneficial results (for instance, in a fully competitive and transparent marketplace). But that may not be the case in consumer markets, to the extent that an imbalance of power exists between those on the supply side and the consumers (on the demand side) with whom they are interfacing.

### Issue: Discrimination in Ad Delivery

Advertisements that give their addressees an advantage by making them aware of good opportunities may be discriminatory against those who are excluded from such communications (e.g., jobs being offered only or mostly to male candidates, houses being provided to people matching the existing ethnic composition, etc.). AI systems that select their addressees in providing such communications may be innocent of intentional discrimination, in the sense that their action is only aimed at maximising successful responses by sending ads and offers only to those who, according to the system's assessment, would most likely be interested in them or only to those with whom there is a suitable match. An assessment of this kind is generally based on the fact that similar persons have taken advantage of such opportunities in the past or that these opportunities in the past were found to be suitable for them. However, this reliance on past practices or assessments may be based on exclusion and prejudice and lead to unfair discrimination<sup>23</sup>.

### Issue: "Turn off" Personalisation

Personalisation has been at the centre of the debate on data and consumer protection. On the one hand, people need help in selecting what is relevant to them out of the huge amount of information and offers which are available online. On the other hand, intermediation, entrusted to platforms and search engines, may lead to outcomes that are suboptimal for users. The intermediaries need to ensure some alignment between what users see and what they like to see (so that they stay on the platform they are on or keep using the service they are paying for) but what is being delivered to users may fail to meet their needs as informed consumers and engaged citizens. Profiled users tend to receive messages that are more likely to trigger desired purchasing or other behaviour<sup>24</sup> or news that is more likely to keep them on the platform, since it confirms their opinions or piques their curiosity. We may ask whether consumers should know what messages are based on profiling of their personality and which aspects were found to be salient in targeting such messages. We may also ask whether consumers should always be offered the opportunity to opt out of profiling, so that they may enter the marketplace as anonymous visitors. An important suggestion in this regard comes from the GDPR, which gives data subjects the right to object to profiling, which in particular includes a right to request termination of profiling for purposes of direct marketing.

### Empowering civil society

To ensure an adequate protection of citizens, besides regulation and public enforcement, the countervailing power of civil society<sup>25</sup> is also needed. Civil society can provide a unique contribution to detect abuses, inform the public, activate enforcement, stimulate and control regulators, etc. Moreover, active citizenship is an important value in itself, that needs to be preserved and advanced at a time in which individuals tend to delegate to technology (and in particular to AI) many of their choices. However, in the AI era, an effective countervailing power needs to be supported by AI too: only if citizens and their organisations are able to use AI to their advantage can they resist and respond to AI-powered companies and government.

A few examples of citizen-empowering technologies are already with us, as in the case of ad-blocking systems as well as more traditional anti-spam software and anti-phishing techniques. Yet, there is a need to move a step forward. On the side of the consumers, services could be deployed with the goal of analysing and summarising massive amounts of product reviews or comparing prices across a multitude of platforms. Such a process would likely build a new market where digital agents are algorithmic consumers, capable to negotiate, form coalitions, and make purchasing decisions.

Machine learning and natural language processing technologies can be used to automatically analyse textual documents in order to validate their content. One example in this direction is offered by CLAUDETTE<sup>26</sup>, an online system for the automatic detection of potentially unfair clauses in online contracts and in privacy policies<sup>27</sup>. Considerable effort has also been devoted to the development of data mining techniques for detecting discrimination<sup>28</sup>, in particular to build supporting tools that could identify prejudice and unfair treatments in decisions that regard consumers. Multiple AI methods to provide privacy and consumer protection could be merged in integrated PDA-CDA (Privacy digital assistants/consumer digital assistants), meant to prevent excessive/unwanted/unlawful collection of personal data as well as to protect from manipulation and fraud. An integrated PDA-CDA could also enable consumers to escape from the “filter bubble” (the unwanted filtering/pushing of information) and provide them with awareness of fake and untrustworthy information.

It may be worth considering how the public could support and incentivise the creation and distribution of AI tools to the benefit of data subjects/consumers. Such tools would provide new opportunities for research, development, and entrepreneurship. They would contribute to reducing unfair and unlawful market behaviours and favouring the development of legal and ethical business models. Finally, consumer-empowering technologies would support the involvement of civil society in monitoring and assessing the behaviour of businesses and of the technologies deployed by the latter, encouraging active citizenship, as a complement to the regulatory and law-enforcement activity of public bodies.

## Conclusion

This briefing has addressed the regulation of artificial intelligence (AI). Focusing on data and consumer protection, it has introduced prospects and risks of the applications of AI, identifying the main legal regimes that are applicable, and a set of key legal issues to be addressed.

It has shown that the many diverse applications of AI provide multiple opportunities and challenges for the European society. It has argued that a regulatory approach is needed, which takes into account fundamental rights and principles as well as existing regulatory regimes, such as data protection, consumer protection and competition law. This would ensure that AI benefits citizens and communities, according to European values and principles. It is also necessary to support the countervailing power of civil society, in particular by incentivising the development of appropriate AI tools.

Consumers are strongly affected by the use AI, in particular with regards to the provision of personalised services, based on the processing of personal data. An adequate regulatory response is needed for urgent issues such as price discrimination, personalised advertising/nudging, discrimination in ads delivery, or the opportunity to turn off personalisation.

- <sup>1</sup> R. Kurzweil, *The Age of Intelligent Machines*, MIT, 1990, 14. For a discussion of different definitions of AI, see S. J. Russell, and P. Norvig, *Artificial Intelligence. A Modern Approach* (3 ed.), Prentice Hall, 2010, Ch. 1.
- <sup>2</sup> Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions: Artificial Intelligence for Europe SWD(2018)137 final.
- <sup>3</sup> L. Floridi, J. Cows, M. Beltrametti, R. Chatila, P. Chazerand, V. Dignum, C. Luetge, R. Madelin, U. Pagallo, F. Rossi, B. Schafer, P. Valcke, and E. Vayena, Ai4people—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations, *Minds and Machines* 28, 2018, 689–707.
- <sup>4</sup> N. Nilsson, *The Quest for Artificial Intelligence*, Cambridge University Press, 2010.
- <sup>5</sup> N. Bostrom, *Superintelligence*, Oxford University Press, 2014.
- <sup>6</sup> A. Halevy, P., Norvig, and F. Pereira, The unreasonable effectiveness of data, *IEEE Intelligent Systems*, 2009, 8–12.
- <sup>7</sup> See N. Cristianini, The road to artificial intelligence: A case of data over theory, *New Scientist*, 2016, 26 October.
- <sup>8</sup> V. Mayer-Schönberger and K. Cukier, *Big Data*, Harcourt, 2013.
- <sup>9</sup> S. Zuboff, Big other: surveillance capitalism and the prospects of an information civilization, *Journal of Information Technology* 30, 2015, 75–89.
- <sup>10</sup> H. R. Varian, Computer mediated transactions, *American Economic Review*, 100, 2010, 1–10.
- <sup>11</sup> C. O'Neil, *Weapons of math destruction: how big data increases inequality and threatens democracy*, Crown Business, 2016.
- <sup>12</sup> E. Pariser, *The Filter Bubble*, Penguin, 2011.
- <sup>13</sup> S. Zuboff, Big other (endnote 9).
- <sup>14</sup> J. M. Balkin, The constitution in the national surveillance state, *Minnesota Law Review* 93, 2008, 1–25.
- <sup>15</sup> Floridi et al, Ai4people (endnote 3).
- <sup>16</sup> S. Wachter, B. Mittelstadt, and L. Floridi (2016). *Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation*.
- <sup>17</sup> F. Pasquale, *The black box society: the secret algorithms that control money and information*, Harvard University Press, 2015.
- <sup>18</sup> On technologies for explaining the function of black-box systems, see R. Guidotti, A. Monreale, F. Turini, D. Pedreschi, and F. Giannotti, A survey of methods for explaining black box models. *arXiv:1802.01933v2 [cs.CV]*, 2018.
- <sup>19</sup> J. M. Balkin, The three laws of robotics in the age of big data, 2017, 1217-241.
- <sup>20</sup> Pariser, *The Filter Bubble* (endnote 12).
- <sup>21</sup> Townley, C., E. Morrison, and K. Yeung, *Big data and personalised price discrimination in EU competition law*. Technical Report 2017-38, King's College London Dickson Poon School of Law. Legal Studies Research Paper Series, 2017.
- <sup>22</sup> K. Yeung, 'Hypernudge': Big data as a mode of regulation by design. *Communication and Society* 20, 2018, 118–36.
- <sup>23</sup> For an overview of the issue of unfairness in algorithms, see S. Barocas, and A. D. Selbst, Big data's disparate impact. *California Law Review*, 104, 2016, 671–732; J. A. Kroll, J. Huey, S. Barocas, E. W. Felten, J. R. Reidenberg, D. G. Robinson, and H. Yu, Accountable algorithms. *University of Pennsylvania Law Review* 165, 2016, 613–705; P. Hacker, Teaching fairness to artificial intelligence: Existing and novel strategies against algorithmic discrimination under EU law. *Common Market Law Review*, 2017.
- <sup>24</sup> This is also called persuasion-profiling, see Pariser, *The Filter Bubble*, 82 (endnote 12).
- <sup>25</sup> J. K. Galbraith, *The Anatomy of Power*, Houghton-Mifflin, 1983.
- <sup>26</sup> The software is available at the link <https://claudette.eui.eu>.
- <sup>27</sup> G. Contissa, F. Lagioia, M. Lippi, P. Palka, H.-W. Micklitz, G. Sartor, and P. Torroni, Towards consumer-empowering artificial intelligence. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*, 2018, 5150–7.
- <sup>28</sup> Ruggeri, S., D. Pedreschi, and F. Turini, Integrating induction and deduction for finding evidence of discrimination. *Artificial Intelligence and Law* 18, 2010, 1–43.

**Disclaimer and copyright.** The opinions expressed in this document are the sole responsibility of the authors and do not necessarily represent the official position of the European Parliament. Reproduction and translation for non-commercial purposes are authorised, provided the source is acknowledged and the European Parliament is given prior notice and sent a copy. © European Union, 2019.

Manuscript completed: December 2018; Date of publication: January 2019

Administrator responsible: Christina RATCLIFF; Editorial assistant: Roberto BIANCHINI

Contact: [Poldep-Economy-Science@ep.europa.eu](mailto:Poldep-Economy-Science@ep.europa.eu)

This document is available on the linternet at: [www.europarl.europa.eu/supporting-analyses](http://www.europarl.europa.eu/supporting-analyses)

IP/A/IMCO/2018-16

Print ISBN 978-92-846-4507-7 | doi: 10.2861/070792 | QA-01-19-044-EN-C

PDF ISBN 978-92-846-4508-4 | doi: 10.2861/441665 | QA-01-19-044-EN-N