PART I

Rise of AI Consumer Markets

1

AI and Consumers

Geraint Howells

1.1 INTRODUCTORY REMARKS

If you skipped to the concluding chapter, which many people might since it is written by a true expert on this topic, Roger Brownsword,[1] you will see that he reflects on the birth of e-commerce and the lack of knowledge we had then of how this field would develop. There were certainly new problems being posed for the law to adapt to e-commerce and to test if new payment methods could be made to work. However, even then my feeling was that these were problems that could be understood within a traditional context. Academics can frequently suffer from an imposter syndrome, as they are given an expert status that they may not be worthy of fulfilling. In truth, most of the time we can add something of value to the discussion. However, in this AI world, my imposter status paranoia is heightened. It was my lack of knowledge of this technology that made the offer by Larry DiMatteo and Christina Poncibò to host a conference in Galway so attractive.

However, I needed more persuading to be an editor of this book, as that suggested I was providing thought leadership in this area. But, as a consumer law scholar, the importance of the intersection between AI and consumer protection law has become obvious. Nonetheless, I have seen too many lawyers expose themselves due to lack of technical savviness and too many tech-savvy people reveal their limited understanding of the law. I have, nevertheless, in the past made some comments on the challenges facing consumer protection in the face of the onward march of technology.[2] The broad thrust of my contributions to date has been that we should be reluctant to throw away hard (and recently) won consumer rights and values on the altar of technological developments. This need to protect such rights was enforced during a prolonged stay in Asia. Consumer law is less well developed there than in Europe[3] and technological development was a far stronger political driver than consumer protection. However, the fast-paced changes in technology in the last few years make me feel that the debate over AI and consumer protection should not become dominated by technology experts. Equally, whilst I still stand by my call to protect traditional consumer protection values (and in many ways the new technologies highlight and emphasise the need for consumer protection), it is becoming obvious

[1] See Chapter 19.
[2] G. Howells, 'Consumer Protection Values in the Fourth Industrial Revolution' (2020) 43 *Journal of Consumer Policy*, 145.
[3] G. Howells, H. Micklitz, A. Janssen and M. Durovic (eds.), *Consumer Protection in Asia* (Hart 2022).

that the future may require at least some changes to the consumer protection landscape to adapt to the new technologies.

My primary role in the creation of this book has been to read the collection of chapters and make some introductory comments. Having read the chapters, one reflection is that there is a new generation of scholars with both the legal and technological skills needed to understand the problems and make valuable contributions to resolving them. I also came away concluding that I am not alone in being concerned about the lack of precise knowledge of the intersection of technology and consumers. We live in a clichéd world of the 'unknown unknowns'. We are operating in a space dense with uncertainties. There is no certainty about what AI means, what it can do, or even what we want it to do. As far as the law's reaction to AI is concerned, there is uncertainty about whether and how traditional law applies or can be adapted or what new approaches should be taken. There is then the question of whether the new approaches reflect new values. Does the law need a new understanding of damages to handle the issues presented by AI? Do we need new regulatory tools to address the new actors? Does AI need discrete legal rules, or can general rules be adequately applied, and can the general and the specific coexist alongside each other?

On reading the chapters, I found at least a dozen areas of uncertainty raised as to the interaction of AI and consumers: (1) uncertainty about what AI is, (2) uncertainty about what AI can do, (3) uncertainty as to what we want AI to do, (4) uncertainty about how traditional rules apply to AI, (5) uncertainty about how far traditional legal rules can be adapted to AI, (6) uncertainty about new approaches to police the dangers of AI, (7) uncertainty about whether AI laws and other laws should overlap, (8) uncertainty about new values underpinning AI regulation, (9) uncertainty about what damages should be recoverable, (10) uncertainty about the roles and responsibilities of new actors, (11) uncertainty about enforcement and (12) uncertainty about the impact of AI on the role of lawyers. The next three subsections discuss these uncertainties in three categories: uncertainty about the application of AI to consumer markets, uncertainty about the legal regulation of AI and uncertainty around enforcement of consumer law in the AI context.

1.2 UNCERTAINTY ABOUT THE APPLICATION OF AI TO CONSUMER PRODUCTS AND MARKETS

Terms such as AI and blockchain are often bandied around without it being clear whether what is being discussed is really related to a given technology. Of course, none of our contributors are ignorant of the true meaning of AI and several chapters seek to explain the technology, though it is noted the definitions remain contested.[4] There is indeed a distinction between AI that makes decisions based on predetermined instructions and more advanced AI that uses machine learning to potentially make decisions in ways that the original programmer could not have predicted. It is these self-learning AI systems around which there is most concern and which potentially pose the biggest challenge to traditional regulation.

These chapters were written at a time when some close to the development of AI were suggesting a pause in its development to bring in proper regulation. The consumer applications are not on the high-risk register – though one could imagine autonomous cars or AI-engineered toys becoming lethal if they developed malicious or simply misguided motives and actions.

[4] See Chapter 14.

There is generally a feeling that AI brings both risks and benefits. The autonomous car, for instance, has the potential to be safer and deliver us faster to our destination than traditional vehicles. Of course, drivers fear the loss of control. It is in the field of advertising and marketing where perhaps the use of AI will have the most impact on consumers. There, the benefits of personalisation must be balanced against the risks, such as discrimination and manipulation.[5] The use of AI can make certain advertising practices even more dangerous when advertisers have the power of AI to enhance them.[6] Indeed the new AI-driven marketing systems bring greater inequality of technology and bargaining power between consumers and traders.[7] Consumers may not be able to recognise the strategies being employed against them or appreciate that the agents being used by businesses are working in the businesses' interests and not theirs.[8] It may be that agents working as substitutes for consumers and working in their interests will become features of the market place, but that has not yet happened.[9] Indeed, we do not yet know the full capacity of AI. For instance, when will the use of visualisation become a feature of recommendation agents.[10]

Beyond knowing what is or will be technologically possible, there is a need to consider what we want to permit. For instance, is personalised pricing to be encouraged or prohibited? Extreme personalisation can remove the category of vulnerable consumer altogether only to reveal how vulnerable we all are.[11] It is also suggested that the acceptance or lack of acceptance of AI may come down to a more general lack of trust in surveillance capitalism that it is associated with.[12] AI technology is moving fast. The answers to what it can do may become clearer, but the debate will then have to be about what we want to allow it to be used for. Technology should be our tool and its deployment needs careful social and technological regulation.

1.3 UNCERTAINTY ABOUT THE LEGAL REGULATION OF AI

One point that was well made in several chapters was that it is difficult to know how to apply existing law to AI.[13] This is because the standards and expectations of AI are not yet settled. The 'blackbox' effect may, for instance, make it impossible to apply the objective criteria of the Digital Contents Directive as it is impossible to know what choices have been made by the AI developer.[14] You cannot assess that which you do not know. It is also noted that the Artificial Intelligence Act uses the phrase 'subliminal' without defining it.[15] Such abstract concepts need to be clearly defined if they are to be operationalised.

There is thus a question over whether the existing rules can even by applied. Though, there is hope that the development of standards and codes will over time flesh out the expectations of AI systems. Of course, that process itself will require a high degree of governance to make the

[5] See Chapters 5, 6 and 11.
[6] Chapter 6.
[7] See Chapter 13.
[8] See Chapter 13.
[9] See Chapter 3.
[10] See Chapter 2.
[11] Chapter 13.
[12] Chapter 11.
[13] Chapter 3.
[14] Chapter 5.
[15] Chapter 13.

outcomes socially acceptable. Running through many of the contributions is the tension over whether existing laws can meet the challenge of AI regulation (adaptions of existing rules) or if more radical thinking is needed. In other words, should we be aiming at coherentism with the existing norms or more instrumental reform.[16]

It is noted that some traditional laws suffer weaknesses that may affect their value in the AI context. Data protection regulation relies on disclosure and consent, but consent fatigue reduces its effectiveness.[17] This is exacerbated in the digital context where tick boxes are used to gain user consent. We need to be alert to the risk of only formal compliance as well as outright evasion of true consent.[18]

Some of the chapters openly suggest that new models may be needed[19] or even new concepts such as the personhood of AI entities.[20] Indeed, technology might itself be the solution to the problems it creates. We may be able to introduce protection by design into AI systems,[21] although it is not clear that such solutions will be widely accepted and adopted by business.

The debate around AI liability is currently ongoing.[22] It questions whether traditional private law liability systems will work or if there is a need for a compensation fund.[23] The question of liability is clouded by the numerous new actors on the scene that may affect who should be responsible or liable. AI systems have both back-end and front-end operators.[24] Platforms may have an important role to play in the regulatory landscape.[25] There is also a fierce discussion about whether the harm suffered and damaged claimed need to be recast for the AI context. At minimum, regulation is needed to protect fundamental rights in the AI context[26] and also to consider whether psychological harm should be covered.[27]

It is noted that the reform proposals in Europe offer good protection for some aspects of AI, but it is limited to the category of high-risk AI systems.[28] The new EU AI Act has an uncertain relationship with EU and national consumer law. At one level, it can be seen as providing *ex ante* controls, whereas consumer law controls what is on the market. Where it does impose substantive consumer rules, it might be questioned if they would be better situated in consumer law.[29] Some regulations, such as the Digital Services Act, only apply to companies that operate platforms and not when they supply directly.[30] Even slight ambiguities in the new legislation will likely be exploited by the well-advised technology companies.[31]

There is clearly a debate to be had about the extent to which AI can be regulated by existing laws or if special laws need to be adopted for this context. For instance, does AI need special regulation in the context, for instance, of data protection?[32] Should AI regulation overlap with

[16] See Chapter 7.
[17] Chapter 13.
[18] Chapter 11.
[19] Chapter 6.
[20] See Chapter 10.
[21] Chapter 11.
[22] Chapter 10.
[23] See Chapter 9.
[24] Chapter 10.
[25] Chapter 13.
[26] Chapter 9.
[27] See Chapter 12.
[28] Chapter 5.
[29] See Chapter 15.
[30] Chapter 13.
[31] Chapter 13.
[32] Chapter 5.

unfair commercial practices law?[33] Should the issue of liability for psychological harm be addressed specifically for AI or more broadly in tort law?[34]

1.4 UNCERTAINTY AROUND ENFORCEMENT OF CONSUMER LAW IN AI CONTEXT

AI also has implications for the enforcement of consumer protection law.[35] It has the capacity to assist in its enforcement and take on tasks (particularly routine tasks) such as detecting websites in breach of consumer protection law. However, an air of caution is needed, as the ability of AI to assist is limited at the moment. This may change as AI systems develop. There are also concerns about opacity (blackbox effect), potential bias, problems of distinguishing correlation from causation and cyber-security risks from the transfer of data.

The scope for AI should include assisting the courts in mass litigation and in appropriate areas like consumer law.[36] So long as AI only assists judges rather than taking over the decision-making process, its use should be compliant with the fundamental right to a fair trial. However, there is a need for caution, as there have been instances of AI tools demonstrating bias. The full potential of its use by the courts is still uncertain.

There are also implications for the work of lawyers as AI becomes more embedded in the practice of law and may assist judicial decision-making.[37] Once again, there is a need to distinguish between current use of AI, known as legal tech, where lawyers have begun to adopt AI systems aimed at improving efficiency and accuracy, and future advanced AI systems that pose dangers for the practice of law. There is the interesting observation that legal tech is more advanced in the common law countries of the United Kingdom and United States. This seems to be partly for reasons such as language and availability of databases but also due to the structure of legal markets in these jurisdictions that are dominated by large firms that can afford to invest in technology. Being based in a small common law jurisdiction (Ireland), it intrigued me as to whether this could be used as a testing ground to discern the factors most important to the deployment of AI. Another common feature between lawyering and AI deployment is the need to fill any gaps in ethical regulation. Lawyers and judges need to develop ethical guidelines, just as society has to work out its general approach to AI governance.

1.5 CONCLUSION

AI is a fast-moving area. Its potential is as yet unknown. The issues it raises are numerous and yet often not yet capable of being clearly formulated. We are uncertain about what it can do, what we want it to do and how we should regulate it. There may be dangers in regulating too quickly as that can be a barrier to innovation. Equally, in time, it may be discovered that the legal system can resolve the issues using existing tools. Courts can be very innovative in finding sensible solutions. Legislators can be relied upon to step in and make technical changes when necessary. However, time may not be on our side and the process of starting the discussion on standards of

[33] Chapter 5.
[34] Chapter 12.
[35] Chapter 14.
[36] See Chapter 18.
[37] See Chapter 17.

conduct expected in relation to AI needs to begin. Society needs to influence and mould our expectations, so AI is used for the collective good. We all need to be engaged in that debate and not just the technologically literate. Reading this collection has given me more confidence to join in the debate. The reader will hopefully have as positive an experience. At the very least, there is some comfort in learning that others share the sense of uncertainty as to how to respond to AI.

2

Artificial Intelligence for Consumers

Advances in Recommendation Systems

Luigi Portinale and Alessandro Abluton

2.1 INTRODUCTION

This chapter explains how new technologies have evolved to aid consumers in navigating complex marketplaces and make choices that best fit their interests. Its focus is on more advanced AI recommender systems (RSs) (suggestion of products to consumers) that have largely been beneficial to modern-day consumers. Artificial Intelligence (AI) is now a mature scientific discipline that pervades numerous aspects of our everyday life. Several researchers have criticized AI as a name for the discipline[1] by stressing the fact that the term "intelligence" has a non-unique meaning, leading to a discipline that has no clear, well-defined goal.[2] However, understanding and replicating human intelligence is an extremely difficult problem and involves sophisticated methodologies, as well as the availability of suitable computational resources and data. The history of AI has thus evolved between ups and downs in the so-called AI seasons, where spring days (with lots of funding and active projects) are followed by winter days (where funding disappears and results disappoint).

A subfield of AI is machine learning (ML); the goal of ML is to provide artificial systems (agents) with learning-from-experience capabilities. Even if ML is often confused with AI, it is a subfield of AI with specific goals concerning automatic learning from data and situations, while AI has a wider objective concerning the construction of an intelligent agent that learns the needed knowledge (through training data), and then uses that knowledge to perform a specific task. For example, old-fashioned *expert systems*[3] were built without any learning component; the knowledge base needed to perform the targeted task (suggestion of a therapy, discovery of a mineral deposit, design of a computer configuration) was constructed through a manual process called knowledge engineering, where the computer scientist had to work with a domain expert, with the goal of transferring the expert's knowledge into the system through the suitable formalism. This approach quickly showed its limits.

Modern AI mainly focuses on ML to develop AI methods and applications. This is due to the availability of three main assets: new methodologies for AI model building and refinement (basic research in computer science), big data of very different types (text, images, sounds, structured

[1] J. Moor, "The Dartmouth College Artificial Intelligence Conference: The Next Fifty Years" (2006) 27 *AI Magazine*, 87–90.
[2] S. Legg and M. Hutter, "A Collection of Definitions of Intelligence" (2007) 157 *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms*, 17–24.
[3] P. Jackson, *Introduction to Expert Systems* (3rd ed. Addison Wesley 1998).

data), and large-scale/high-performance computational resources. The rise of the so-called deep learning (DL) approach[4] allows us to address very difficult problems requiring complex reasoning capabilities as in decision support, reactive behavior (autonomous devices), and interactive behavior as in personal assistant devices understanding natural language. DL involves training computer systems, called neural networks, to recognize and interpret information in a similar way to how our brains process data. The term "deep learning" refers to the fact that such networks have a very long or deep structure, with several layers devoted to the interpretation of the input in terms of abstract features. For example, in interpreting the picture of a person's face, the initial layers would be devoted to recognizing edges, deeper layers may recognize specific parts of the face (eyes, nose, mouth), and even deeper layers may recognize a specific category of the face (child, male or female).

The impact of such methodologies pushes the current technologies toward new applications by producing complex side effects on the everyday environment, in particular at the socio-economic, juridical, and ethical levels. The study and exploitation of consumer behavior is not an exception. Consumers are a source of a great amount of data involving their purchases, preferences, and intentions. User profiling for improving revenues in e-commerce applications has become a common activity for suppliers of products and services. To this end, a new category of machine learning applications called recommender systems has evolved. A recommender system is a machine learning system that can recommend to a given user new items (products or services) that she would probably like.

This chapter will review the mainstream approaches that have been developed for RSs, with particular attention to the difference between content-based and collaborative filtering approaches. It will discuss current trends in RSs, such as the production of visual recommendations that exploit images in the process of identifying relevant items for the given user-consumer. We review the basic approaches to the construction of RSs, discuss the problem of recommending new products to consumers through images, and report on some specific functionalities that are implemented in a commercial product called *Visidea* developed by Inferendo.[5]

2.2 RECOMMENDER SYSTEMS: TRADITIONAL PARADIGMS

Item recommendation approaches can be divided into two broad categories:[6] personalized and non-personalized. Non-personalized recommendations are the simplest form of suggestion, such as suggesting popular items in a time frame or in a particular geographical region. These recommendations are made to everyone in a particular group of users. Of course, using generic information as in the case of non-personalized recommendations has obvious shortcomings, since recommending items to a customer is an activity requiring knowledge about her preferences in order to obtain reasonable results.

Personalized approaches, on the other hand, try to exploit a wide set of data about users and their interactions within online markets to provide user-tailored recommendations by leveraging machine learning algorithms. In order to perform such computational tasks, RSs collect information from users regarding their preferences, which are either expressed explicitly (ratings for products) or inferred by interpreting the actions of the users. For instance, an RS may consider the navigation to a particular product page as an implicit sign of preference for the items shown

[4] I. J. Godfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press 2016).
[5] See www.inferendo.ai.
[6] F. Ricci, L. Rokach, and B. Shapira, *Recommender Systems Handbook* (Springer 2010).

on that page. Traditional RS paradigms can be organized by means of a taxonomy provided by Robin Burke[7] that identifies the following categories: content-based, collaborative filtering, demographic, knowledge-based, and hybrid systems. Demographic systems make suggestions based on the demographic profile of the user, while knowledge-based systems recommend items based on specific domain knowledge (explicitly modeled through specific formalisms such as knowledge graphs and rules) about the importance and behavior of specific items' features. The most widely used approaches are content-based, collaborative filtering, and hybrid systems, which are discussed in Sections 2.2.1–2.2.3.

2.2.1 *Content-Based Recommenders*

The first approach to RSs is based on a content-based paradigm. When using a content-based approach, the system learns to recommend items most similar to the ones that the user bought or liked in the past. Similarity between items is a key concept in this approach, and it is based on an explicit description of each item in terms of given features. Similarity is then calculated by comparing the features associated with the considered items. For example, if the user showed a preference for watching comedy movies, then the system will learn to recommend other movies from the same genre. Content-based (or cognitive) methods exploit data relating to the user's purchases and interests in an effort to identify the common characteristics of items that have received favorable ratings from the user and then recommend new items that share these characteristics.

RSs built with content-based techniques are generally simpler to develop than other approaches (such as collaborative filtering) but suffer from several problems. First of all, a specific and detailed description of the items of interest (and the users as well) must be available to the system. Moreover, they encounter problems of limited content analysis and overspecialization. Limited content analysis issues arise when a system does not have enough information on the users' preferences or the common characteristics of the items of interest. The solution is to gather more data on the characteristics of the targeted items and greater tracking of the users' actions. The overspecialization or diversity problem is another common issue in content-based RSs and concerns the fact that recommendation algorithms suggest goods directly related to the customer profile rather than new types of items. To depict such a problem, consider an RS that continues to suggest movies of the same genre and with the same actors, thus failing to recommend items that are different but still interesting. It is important to keep some degree of diversity in recommendations, and to only consider item features found in a user's history prevents broader recommendations. This problem relates to a more sophisticated issue related to the notion of novelty and diversity or what is called the lack of serendipity. By "serendipity" we mean the occurrence of an unplanned but interesting discovery. This feature is considered an added value for RSs, since the recommendation of something that is interesting but unexpected is usually perceived as an *intelligent* suggestion.[8] In current practice, the majority of RSs strive to introduce novel and diverse suggestions that may pleasantly surprise customers, but it is difficult for purely content-based systems to achieve.

[7] R. Burke, "Hybrid Recommender Systems: Survey and Experiments"(2002) 12 *User Modeling and User-Adapted Interactions*, 331–370.

[8] R. J. Ziarani and R. Ravanmehr, "Serendipity in Recommender Systems: A Systematic Literature Review" (2021) 36 *Journal of Computer Science and Technology*, 375–396.