Regulating Recommending: Motivations, Considerations, and Principles

Jennifer Cobbe and Jatinder Singh [1]

Abstract

Internet regulation and 'online harms' are matters of much political and regulatory attention. This debate is beset by issues, including defining 'online harms', respecting freedom of expression, and others. While much of this debate has focused on content hosted by online platforms, comparatively little attention has been paid to the central role of algorithmic personalisation - or ' recommending ' - by platforms in content dissemination in online environments and the problems to which this contributes. Focusing on recommender systems, i.e. the mechanism by which content is recommended by platforms, provides an alternative regulatory approach that avoids many of the pitfalls with addressing the hosting of content itself. This paper therefore explores motivations and considerations for regulating the use of recommender systems by online platforms. In doing so, this paper establishes a typology of online recommending, sets out various problems and consequences of recommending, and argues that recommending content is not one of the three activities for which information society service providers are afforded liability protections under the E-Commerce Directive. To address the identified problems and fill this legal gap, this paper proposes some principles for future regulation, and discusses approaches to oversight and compliance that could work with these principles.

Keywords: internet regulation ; platform responsibility; intermediary liability; E-Commerce Directive; recommender systems

1. Introduction

The past two decades have seen the emergence and growth of digital platforms, [2] positioning themselves at the centre of 'multi-sided markets' [3]; intermediaries between individuals, businesses, organisations, governments, politicians, the public, and others. These platforms often adopt surveillance business models, whereby user behaviour is tracked and analysed to predict future behaviours and interests, personalise services, facilitate behaviourally-targeted advertising, and grow user engagement, platform revenue, and market position.[4] In law, these platforms are typically considered to be providers of information society services [5] ('service providers' [6]).

Online platforms have been associated by some with a variety of 'online harms' [7] (although this term has proven predictably difficult to define [8]). Various regulatory responses to these have been proposed and discussed. [9] These typically focus primarily on controlling or limiting the hosting of content believed to be in some way harmful, [10] and would in practice involve filtering, moderation, and takedowns.

However, this is difficult if not impossible to properly and fully automate, [11] and there are serious concerns for the psychological well-being of human moderators exposed to large quantities of illegal or potentially harmful content. [12] As well as this, freedom of expression issues arise from regulation. The internet plays a key role in contemporary society as a communications medium, and interventions focused on transmitting or hosting content are in fact focused not on platforms themselves but on the communications of individuals. Ultimately, providing a means for communication will inevitably result in some people communicating abuse, incitement, extremism, defamatory material, IP-infringing material, misinformation, disinformation, and so on. All of these could be argued to be potentially harmful in some way, but only some are unlawful. The need to protect freedom of expression in a democratic society means that the risk of individuals communicating undesirable but lawful content must to some extent be accepted in order to protect the fundamental right of freedom of expression.

However, when it comes to systemic societal issues like disinformation, conspiracy theories, violent extremism, and political manipulation, content is not by itself the problem. On its own, or viewed by only a small audience, a video promoting a conspiracy theory is not a public policy issue. It becomes one when it has a large audience, and when it combines with other, related content that works to reinforce the message. Where content is algorithmically disseminated through recommending, this (a) increases its audience, potentially significantly, and (b) typically puts it alongside other, similar content. Rather than speaking of 'harmful' content, then, it is perhaps more accurate to talk about 'potentially problematic' content. That is, content that by itself or when seen only by a relatively small number of people is not necessarily an issue, but when algorithmically combined with other, similar content or disseminated to a large audience can contribute to systemic problems. Interventions focused on the hosting of content itself miss, to a large extent, issues relating to algorithmic dissemination.

This paper therefore argues that regulatory interventions should instead be directed towards the recommending of content by platforms. For issues which operate on a more individual level, such as harassment, abuse, and IP infringement, the content itself is often the problem. But for those which exist on a more systemic plane, such as in surveillance business models, platform monopolisation, voter manipulation, disinformation, and the promotion of violent extremism, recommending plays a significant role.

Fundamentally, recommending by platforms is largely concerned with showing people whatever the platform's algorithms predict will drive engagement, revenue, and market position, [13] often with little concern for what the material being disseminated actually is. As a result, dissemination by recommender systems - the algorithmic mechanism by which content is recommended to users - can amplify issues caused by potentially problematic content and transform content which by itself may be relatively innocuous into a more serious issue. Moreover, the same technical systems also provide the mechanism for the delivery of behaviourally-targeted advertising (a form of paid-for recommendation that can broadly be considered 'content' in a more general sense). Recommending is therefore important for the growth and dominance of monopoly-like platforms, [14] and is at the heart of surveillance business models. [15]

Many platforms have dominant positions in their particular market, and use recommender systems to disseminate content, determining what and how content is recommended, thus giving them great power and influence. This should come with certain responsibilities and, potentially, with certain liabilities. A focus on recommending would allow the development of legal responses acknowledging that, in recommending content, platforms are operating beyond the limits of where liability protections are provided to intermediaries under the E-Commerce Directive. It would also focus attention not just on platform dominance or the hosting of content itself, but on how recommender systems are designed and used by platforms to disseminate content, build revenue and market position, and leverage the power that market position provides. Indeed, in February 2019, the Council of Europe recognised that attention

should be paid to the significant power that these systems confer on those who operate them and encouraged Member States to take steps to address problems this may cause. [16] And in March 2019, the Communications Committee of the UK Parliament's House of Lords, in the report published following its inquiry on internet regulation, found that the E-Commerce Directive was inadequate to govern algorithmic content dissemination and that reform or replacement was needed. [17]

Focusing on the recommending by service provides of user-generated content, this paper intends to establish the motivations for regulating recommending and propose some key principles to inform future law. In doing so, this paper describes recommending and undertakes a review of its use on various platforms and establishes a typology, sets out some of the motivations for regulating recommending, identifies a significant gap in the current liability protection regime applying to online intermediaries in relation to recommending, and proposes some high-level principles to be taken into account in developing regulation in this area. Potential consequences and limitations of this regulatory approach will then be set out, and directions for future legal and technical research identified. In all, this paper proposes the fundamentals of a framework that seeks to align incentives for platforms to proactively work towards responsible recommending.

2. Definition, Examples, and Typology

Recommending involves the algorithmic selection by service providers of 'content' served to individuals or groups according to some determination made by the service provider of relevance, interest, importance, popularity, and so on to those individuals or groups. Providing a search function is not recommending (unless ranking of results is influenced by predicted relevance, interest, and so on to the user rather than, for example, solely by relation to the search criteria itself provided by the user) [18] ; nor is providing an ability to browse content, providing a directory of content, or allowing people to share links to content (whether shared on the platform itself or elsewhere). Providing a chronological feed of content posted by accounts or content providers followed by a user is similarly not recommending. Recommending includes, among other things, constructing feeds of content such as Facebook's News Feed or Twitter's Timeline algorithmically rather than chronologically, determining the ranking of content based on criteria that includes the feedback of other users, and providing suggested or promoted content to users.

Recommending involves the use of recommender systems. Algorithmic systems - of which recommender systems are one form - are not neutral. [19] Any given algorithm exists because somebody somewhere has an outcome that they wish to achieve through algorithmic mediation, whether that's personalising content, encouraging continued platform engagement, delivering behaviourally-targeted advertising, or something else entirely. As Beer says, "Algorithms are inevitably modelled on visions of the social world, and with outcomes in mind, outcomes influenced by commercial or other interests and agendas". [20] Indeed, at a high level any algorithm can be understood to consist of a sequence of steps intended to produce a desired outcome. [21] Algorithmic systems generally and recommender systems specifically are therefore inherently normative in nature. They are also contextual and contingent in nature, in that they are always embedded within and a product of the wider socio-technical context of their development and deployment [22] ; not just the goals of the organisation in question, but also the assumptions, priorities, and practices adopted by engineers, designers, managers, and users. As a consequence of their normative, contextual, and contingent nature, their use can never be neutral.

THE BUSINESS OF RECOMMENDING

Recommender systems are widespread across the internet and play a key role in surveillance business models. These business models (which have been collectively termed 'surveillance capitalism') [23] involve the surveillance and modification of human behaviour for profit, [24] and drive extensive data gathering and analysis, producing serious privacy and data protection concerns. [25] Platforms obtain as much data as possible about as many people as possible doing as many things as possible from as many sources as possible. These datasets are then algorithmically analysed so as to spot patterns from which interests, preferences, and future behaviours can be predicted. The more data that can be gathered about people and their interests, preferences, and behaviours, the more accurate (in theory) those predictions can be.

Recommender systems play two fundamental roles in surveillance capitalism. The first of these is in delivering behaviourally-targeted advertising and other paid-for content to bring direct revenue from advertisers and others. The second is personalisation to drive engagement, thus indirectly contributing to the maintenance of direct revenue streams. While behavioural advertising typically involves building profiles of users, recommending for personalisation does not necessarily. But they do both involve the selective dissemination of content to audiences in pursuit of the platform's business goals and according to the platform's own analysis. That is to say, showing selected content to users that the platform has determined might modify their predicted behaviour in some way - either to persuade them to click on advertising or other paid-for content, or to persuade them to stay engaged with the platform. This paper focuses on recommending by service providers to personalise services and drive engagement, profit, and market position, rather than on behavioural advertising.

EXAMPLES

As Table 1 indicates, Recommending plays a central role across the most popular websites and platforms on the internet.

Organisation	Primary Function	Rank	Description
Google	Search	1	By default, search results are personalised according to Google's determination of the user's interests, alongside information relating to the search query, similar queries by other users, links to related news stories, etc.
YouTube (Google)	Video streaming	2	The home page recommends videos based on recent uploads, popularity, etc. For logged in users, videos from subscribed channels are also recommended as well as videos and channels based on interests and viewing history. Alongside playing videos, links to recommended videos based on similarity to the playing video and to viewing history are provided. After each video finishes, further recommended videos are displayed (a recommended video will, by default, play automatically after a short period of time). Feedback options are provided in the form of 'like' and 'dislike' buttons.
Facebook	Social media	3	News Feed by default provides an algorithmically-produced feed of content of various kinds (posts, likes, comments, etc.) determined to be most interesting or relevant to the user. Showing 'most recent' content is optional (though chronological ordering is the default when viewing an individual's profile). Alongside News Feed, Facebook also displays information on 'Pages' similar to those liked by the user, people who may be known by the user but who have not been added as a 'Friend', and assorted other recommended content. 'Reactions' (including 'Likes') are available as feedback options.
Yahoo!	Search	9	Search results are personalised based on user's search history, and a home page feed of news and other content is recommended based on popularity and the user's predicted interests.

Amazon	Commerce	10	The home page recommends items and categories of items, including 'featured' recommendations. The 'Your Amazon' page similarly recommends items grouped by category. On product pages, similar items are listed, as are items viewed by other users and information on which items other users ultimately purchase after viewing the product. Again, 'featured' recommendations are listed at the bottom of the product page. Upon adding products to the basket, sponsored and similar items are recommended
Twitter	Social media	11	By default, tweets and retweets by followed accounts are interspersed with a selection of recommended content (likes, replies, and so on). Recommended 'top tweets' are shown at the top of the user's timeline. 'Trending' topics are shown alongside the timeline, as are suggested accounts to follow. Users can switch from this algorithmic timeline to a chronological timeline.
Instagram (Facebook)	Social media	16	Posts (including adverts) are ranked algorithmically with no option for a chronological display. Suggested users to follow are provided alongside the content feed. The 'explore' page displays a list of suggested users, recommended content, and adverts.
Reddit	Social media	21	The 'frontpage' by default recommends selected 'hot' posts from across the site. If a user is logged in, this is shows the 'hot' selection from the user's subscribed subreddits. Posts can also be ranked by 'new' (i.e. chronologically) or by 'controversial', 'top', and 'rising' (i.e. algorithmically). 'Upvote' and 'downvote' buttons provide feedback options. A similar arrangement in terms of both ranking and feedback exists for comments within posts.
Netflix	Video streaming	24	The home page recommends videos (including 'top picks' for the user) and categories of videos. At the end of each video further videos to watch are recommended. Each video is assigned a personalised 'match' percentage indicating the user's predicted interest. Users can provide feedback on each video in the form of 'thumbs up' and 'thumbs down' buttons.

Table 1: *An indicative survey of some of the top 30 most visited websites.* [26]

As the Table indicates, the normative power of recommender systems is central to their use in online spaces for benefit of the platforms themselves. They are employed by platforms for a variety of reasons determined by platforms to pursue a variety of outcomes desired by platforms and thus play a key role across large parts of the online ecosystem. On Google, recommender systems are used for various purposes, including to personalise search results to show links that bring revenue to Google. On sites like YouTube, Facebook, Reddit, Twitter, or Instagram, recommender systems provide a personalised feed of content for each user so as to keep them engaging with that platform and drive advertising revenue. Netflix uses recommender systems to present a personalised video library to users as well as recommendations for further viewing so as to keep users watching and subscribing. Amazon uses a recommender system to respond to predicted user desires in order to induce them to buy products from Amazon rather than elsewhere. More recently, popular services on which content is delivered almost exclusively by recommender systems - such as TikTok - have emerged.

TYPOLOGY

At a technical level, the two most common approaches to recommending are content-based filtering and collaborative filtering. [27] Both use machine learning, which produces statistical models trained on (usually large) datasets and can spot correlations and patterns from which to make predictions and draw inferences. [28] Content-based filtering systems recommend content based on similarity to content previously consumed by the user (for example, 'picture X has a similar title to previously viewed pictures Y and Z'). Collaborative filtering systems recommend content based on what similar users have consumed (for example, 'people A, B, and C like this; a similar person D might also like this'). Both involve filtering content so as to show to users only that which is determined by the platform to be relevant, appropriate,

interesting, and so on. Some platforms make use of hybrid approaches, combining features of both methods of filtering. [29]

It is important to not predicate regulation on any particular technical approaches, as new mechanisms may emerge in future. Instead, addressing recommending as an activity allows the law to establish standards that apply to recommending undertaken by any form of system, whether currently used - indeed whether currently existing - or not. To that end, it is useful to distinguish different approaches to recommending in such a way as to develop a typology which stands independent of the technical system used by any particular service. At this non-technical level there are essentially three forms of recommending:

Open recommending. This provides recommendations from a pool of content which is primarily user-generated or submitted, brought in automatically from various sources, or otherwise aggregated in some way without being specifically selected by the platform (although platforms may include other sources of content and may include their own content alongside that generated by users). Google, YouTube, Facebook, Reddit, Instagram, and Amazon all make use of recommender systems in an open manner. On YouTube, for example, any video uploaded by users is normally by default brought into the recommender system.

Curated recommending. This differs from open recommending in that the system selects from a pool of content which is curated, approved, or otherwise chosen by the platform rather than provided directly by users or advertisers or automatically brought in from elsewhere. Netflix is a popular example of a curated system, in that the videos in its library are selected by Netflix (others include BBC iPlayer, Spotify, Apple Music, and so on). Like open systems, curated setups are predominantly adopted where platforms do not produce their own content or where they blend their content with that created by others. Unlike open systems, however, curated systems do not typically include user-generated content without some kind of editorial process. They are often used where more traditional forms of media requiring licensing are involved, such as music, films, or TV shows.

Closed recommending. This is where the content to be recommended is generated by the platform itself or the organisation which operates that site. For example, where a news organisation provides a personalised feed of stories and articles to its users, all of which are produced or commissioned by the organisation itself.

The distinction lies in the platform's role in the sourcing of content it recommends . If the system selects only from the platform's own content, then it is a closed system. If it selects only from content that has been chosen or licensed by the platform (possibly but not necessarily including material produced by the platform) then it is a curated system. If it primarily selects from content provided by users without the platform editorially reviewing that content, then it is open (even where such a system also includes content produced by the platform itself [30]). A system where the default is to include user-generated content in the recommender's source pool but where certain users or certain items can be excluded following terms of service violations, for example, is an open system. Adopting this non-technical typology allows recommending as an activity to be considered in a technology-neutral and platform-neutral way.

This paper is primarily concerned with recommending for disseminating user-generated content or for determining the information to be shown to users without editorial selection; that is to say, open systems, potentially adopting any technical approach to recommending. Of the three forms of recommending, open recommending is the biggest contributor to systemic issues. While closed and curated systems could also be problematic, certain distinctions between those and open systems mean that the latter are more likely to relate to individual-level issues [31] and to contribute to the development of systemic societal issues online. It is possible that a curated system could recommend

hate speech or violent extremism, for example, but the editorial control that distinguishes these systems from open systems means that in reality this is significantly less likely than with open systems. Similarly, it is certainly also possible that a curated system could disseminate IP-infringing content, for example, but the contexts in which these systems are currently used mean that in practice this is unlikely.

3. Recommending and Systemic Societal Issues

While recommending is often touted as a means of personalising services for users' benefit, it is ultimately undertaken to serve the interests of the platforms themselves - to encourage users to stay engaged with the platform in question (whether to consume content, to provide content, to make a purchase, and so on), to bring revenue, and to build market position. [32] Seaver's study of the use of recommender systems shows that they are often intended to 'hook' people, leading to the conclusion that they are essentially employed as 'traps'. [33] Other work produced by YouTube [34] , Netflix [35] , and Yahoo and the website Etsy [36] confirms the view that recommending is fundamentally about engagement in pursuit of profit. As a result, around 80 per cent of Netflix viewing hours come through recommendations; around 20 per cent from its search function. [37] Around 70 per cent of YouTube video views similarly come from recommendations. [38] Understood in this way, recommending is not so much about personalisation, as about showing people the content that the platform predicts will result in the greatest engagement. That is to say, rather than showing people what they want to see, recommending shows people what the platform wants them to see.

Recommending contributes to the increasingly monopolistic nature of dominant platforms through personalisation and driving engagement, amplifying already influential network effects [39] . The dominance of those platforms gives them significant regulatory power. [40] This paper will argue that this power extends to, among other things, constructing online spaces for profit; leveraging their dominance to effect changes in adjacent markets and elsewhere; and engaging in rule-setting on their own platforms. This power is itself primarily exercised through recommending. Recommender systems are therefore both drivers of platform dominance and levers of the power that dominance produces. These systems play a key role for platforms in enforcing norms and influencing behaviour, a situation described variously as 'governance by software' [41] ; 'algorithmic regulation', [42] and 'algorithmic governance' or 'algocracy', [43] and as a form of 'algorithmic governmentality' [44] . As a consequence of their central role as drivers of platform dominance and levers of platform power, and providing the motivation for regulation, recommender systems play a major role in various systemic problems which have developed online.

ALGORITHMICALLY CONSTRUCTING ONLINE SPACE

The effect of widespread recommending by dominant platforms is that much online space is constructed by algorithmic systems. Given the increasingly central role of these spaces in contemporary everyday life and in shared social experience, it could be said that subjective reality as experienced online is itself increasingly constructed by algorithm. [45] While, of course, television and the mass media have long contributed to collective understanding of the world, this was never as personalised, never as involved in mediating communications between individuals, and never as embedded so deeply in constructing the everyday reality of millions of people. Further, it was always possible for society in general to collectively 'see' what television networks were showing, which is currently not the case with the ever-changing, highly personalised outputs of recommender systems, particularly when disseminating user-generated content. Indeed, television and the mass media can be analogised with curated and closed recommending, in that television networks and media outlets exercise control over the content that they broadcast or published - there is no offline equivalent of open recommending in these forms of media.

Recommending's contribution to the dissemination and amplification of disinformation, extremism, and other potentially problematic content results from its role in showing people what the platform wants them to see (in order to drive engagement and so on, as previously described). The personalised and dynamic nature of online space, produced by recommender systems, allows platforms to systematically present users with choice architectures that, as Susser et al put it, "can be specifically designed to exploit each individual user's particular vulnerabilities, and can change and adapt over time". [46] As discussed previously, there are two ways this plays into profit-making: directly through behavioural targeting and (the focus of this analysis) indirectly through personalisation to drive platform engagement. In most cases, the recommending involved in this is open in nature, whereby user content is posted or uploaded and available to be recommended by the system with little concern for the nature of the content itself.

Recommending is now, as Gillespie argues, "a key logic governing the flows of information on which we depend". [47] The choice of what to display and what to not display, exercised by corporate algorithm, could have a significant effect on collective awareness of politics, current affairs, and scientific consensus. For example, as Tufekci points out, [48] the protests and subsequent unrest which erupted in Ferguson, Missouri following the killing by police of Michael Brown in 2014 dominated Twitter discussion in the United States (at the time, Twitter still operated a chronological timeline) and from there entered into the mainstream media. Yet stories about the protests or the resulting Black Lives Matter movement barely surfaced on Facebook's algorithmic News Feed for quite some time because content relating to the 'ice bucket challenge' was receiving more 'Likes'. Tufekci calls this 'algorithmic gatekeeping' - "the process by which ... non-transparent algorithmic computational-tools dynamically filter, highlight, suppress, or otherwise play an editorial role - fully or partially - in determining information flows through online platforms and similar media". [49]

This gives platforms great influence over the shaping of online spaces. The algorithmic construction of social media and other services and the resulting corporate lens through which reality is viewed means that users across the internet often do not get a true picture of what's going on in the world. Through the construction of online space, platforms can influence users' subjective understanding of the world and of their own experiences. [50] Indeed, Facebook's own published research showed that it can actively influence users' emotional state by tweaking its algorithm to show more positive or negative content. [51] The construction of online spaces is itself a form of corporate algorithmic influence, a softer form of behavioural hypernudging [52] facilitated by platform dominance and driven by desires for engagement, revenue, and market position.

Dissemination and amplification: Systemic issues

In many cases, systemic issues do not arise through content alone; dissemination by algorithm plays a key role in amplifying audiences, thereby taking individual content items and producing systemic societal consequences. Recommender systems often also place content alongside other content of a similar topic or nature, potentially further contributing to the development of problems at a systemic level. It arguably doesn't particularly matter if the content is intended to spread disinformation [53] or a conspiracy theory is seen by a small number of people. The content itself is not necessarily a problem. But if the audience for that content is algorithmically amplified by a recommender system (for example, by recommending that viewers of other, perhaps more truthful, content watch it next) then that could be an issue. And if that content is placed alongside other, similar items, this could effectively modify the choice architecture to direct users' attention and selections increasingly towards content of that nature. The more that users interact with that content - whether by viewing it, liking it, or sharing it - the more likely it is to be recommended in the same way, which would potentially result in more viewing, liking, and sharing, thereby leading the recommender system to promote it further. If, as a result of dissemination through such algorithmic feedback loops, the disinformation reaches a large audience, then content which is in and of itself quite harmless could potentially have become part of a systemic problem.

Recommending thus contributes to the development of various problems, which are themselves amplified by monopolisation. Prioritising for engagement is likely to favour content that produces an emotional response and therefore may be controversial, shocking, or extreme, as people tend to be drawn to this content. [54] Indeed, as a result, academic research and journalistic investigations have shown that recommending can play a significant role in the spread of material promoting violent extremism, neo-Nazism, and Holocaust denial. [55] In particular, recommending seems to have contributed to the growth of the white supremacist and increasingly violent [56] 'alt-right' [57] and, as part of that, in the development of 'GamerGate' [58] (a misogynist harassment campaign masquerading as a debate about 'ethics in video games journalism' which spread to several platforms as well as the mainstream media [59] , helped grow and embolden the contemporary 'alt-right' [60] , and has had a significant impact on public discourse more generally in the US [61]). Similarly, recommending has repeatedly been found to play a major role in spreading disinformation and conspiracy theories. [62] There is also concern about platforms recommending content promoting self-harm, eating disorders, and anorexia to vulnerable users. [63] It appears that seeding 'counter-messages' (content intended to counter disinformation, conspiracy theories, and so on) may merely result in the very content they are intended to 'counter' being recommended alongside them, thereby potentially increasing exposure. [64] It could then become impossible for those seeking to counter systemic problems to use platforms like YouTube to do so effectively, while those seeking to spread disinformation or conspiracy theories continue to proliferate. Indeed, a pro-vaccination charity announced in early 2019 that it had been forced off YouTube as a result of this effect. [65]

Further, platforms not only amplify messages through algorithmic dissemination but can also influence the content being produced. Recommending's amplification effect, coupled with the consolidation of power in online platforms, works to incentivise the production of certain kinds of content (for example, that most likely to 'trend'), thereby increasing the prevalence of particular viewpoints. That is, content that gains a larger audience on a platform, driven in large part by recommending, naturally encourages and incentives the production of similar and related content, by the original producer and by others. This may be exacerbated by the platform's monetisation programme, which can financially incentivise the creation of material that is tabloid, controversial, or otherwise produces an emotional response (which, as previously noted, is often prioritised by recommender systems optimised for engagement). Savvy content producers, seeking notoriety, monetary gain, or simply to push a particular message to a broader audience, could tailor their work so as to increase its dissemination by the recommender systems of the platforms they target. The recommending of this content may 'snowball', through feedback loops created by other users responding to - either positively or to challenge - the content they are exposed to, by commenting, sharing, and so forth.

Related are concerns about recommending leading to fragmentation through 'filter bubbles' [66] (by narrowing the range of content recommended to users), echo chambers [67] (by recommending content that reinforces users' interests), and polarisation [68] (by recommending content of an increasingly extreme nature). This seems to have not yet occurred on personalised news services. [69] However, fragmentation has been repeatedly observed on several social media platforms to varying degrees. [70] While this is fundamentally a result of recommending reinforcing the existing psychological biases of users, [71] it appears that the design of recommender systems influences the extent to which this occurs. [72] That said, across this research, the fragmentary effect of recommending seems to be relatively small. And, as Borgesius et al note, studies on polarisation which focus on political systems dominated by two parties, such as in the UK or the US, may not translate easily to other countries with more diverse political landscapes. [73]

The metrics that underpin recommending can be manipulated by automated accounts ('bots') in order to further the spread of potentially problematic content, an activity that Leiser calls 'cyberturfing'. [74] Bots

attempt to 'game' recommender systems by inflating the 'reputation' of content and thus increase its likelihood of being recommended or its position in algorithmic content rankings. By strategically posting content, and artificially inflating views, likes, shares, and other metrics, networks of bots can together shape the construction of online spaces and thus of the content displayed. [75] This is usually undertaken for commercial or political purposes. [76] The widespread use of bots in relation to political content, in particular, has been repeatedly observed across electoral cycles in multiple countries. [77] Bots can artificially seed political messages in organic discussions, bring greater attention to stories (real or fake), and boost ideas (fringe or otherwise) into mainstream discussion.

Collectively, these studies and investigations show that open recommending can play a significant role in the dissemination of disinformation, conspiracy theories, extremism, and other potentially problematic content. Generally, these recommender systems do not deliberately seek to promote such content specifically, but they do deliberately seek to promote content that could result in users engaging with the platform without concern for what that content might be. Recommending does not in and of itself cause these problems - their roots often lie in social, political, and economic causes, or just in basic human nature. But the prioritisation of engagement in recommending is key to exacerbating those issues online . A system primed for engagement that realises that people are prone to 'rubbernecking', for example, might conclude that everything should be car crashes. As a result of prioritising engagement, the content that is most sensational, most dramatic, most controversial, and most attention-grabbing becomes prioritised at the expense of things that are more mundane (which could, in turn, influence the kinds of content that is produced). Theoretical content neutrality therefore has the practical effect of amplifying potentially problematic content ahead of other material. Recommending by platforms in pursuit of engagement and profit without considering these effects and without restrictions can compound those underlying issues by amplifying the audience for content which is potentially problematic at a systemic level, making it easier to find other, similar content, and facilitating manipulation of the recommender systems themselves.

LEVERAGING PLATFORM POWER

Platforms can use recommending to leverage their regulatory power in several ways. The most obvious is through changes to the recommender algorithm to promote or demote certain kinds of content, inducing those using the platform to modify their behaviour in some way desired by the platform. For example, Facebook has repeatedly altered its News Feed algorithm to reduce 'organic reach' (the number of people likely to see content posted for free by Pages). [78] By changing their recommender algorithms, they constructed the environment in such a way that organisations are left with little choice but to pay Facebook for recommendations in the form of targeted advertising if they want to reach wider audiences. [79]

Recommending can also be used by platforms to leverage their dominant position to effect changes in adjacent markets and elsewhere. Changes made to recommender systems can have a significant effect on markets that have come to rely on those platforms in some way. [80] For example, Facebook's announcement in 2014 that it would tweak its News Feed recommender algorithm to prioritise video content heralded a 'pivot to video' across news media as publishers attempted to ensure that their content would maintain its position on News Feed. [81] Facebook's 2016 admission that it had overestimated engagement metrics for video content [82] led to a subsequent algorithm change in 2018 intended to reduce the prominence of video in favour of 'meaningful interactions'. [83] While this does seem to have significantly increased engagement, [84] as Facebook hoped, it also appears to have contributed to a wave of news media redundancies and closures in late 2018 and early 2019 [85] and to an increase in news content relating to divisive political topics and increased use of the 'angry' reaction

button. [86] Through changes to their recommender algorithm, Facebook has had a significant - and significantly negative - effect on news journalism and on discourse more generally.

In addition to influencing the nature of the content being produced or consumed, the ability to affect adjacent markets also manifests in other ways. For example, the European Commission fined Google €2.4bn for systematically promoting its own Google Shopping service and disadvantaging competitors. [87] In January 2019, France's competition regulator opened an investigation into Google's search rankings [88] (which are personalised, rather than simply being ranked according to similarity with the search term). And in 2010, Google announced that it had tweaked its search result ranking algorithm to reduce the prominence of websites that "in [their] opinion" provide a poor user experience. [89] Introducing such subjective elements into recommending establishes a standard for the business practices of others, attempting to influence the behaviour of actors in adjacent markets.

The desire for data about user behaviour in order to 'improve' recommender systems (in the sense of making them better at driving engagement) has also been a motivating factor in platform acquisitions of emerging competitors and actors in other markets. [90] Some responses to this have emerged - the German competition authority has recognised Facebook as the dominant company in social networking and prohibited it from combining WhatsApp and Instagram data with Facebook user accounts without explicit user consent out of concern for abuse of market power. [91] In 2017, Facebook was fined €110m by the European Commission for falsely telling the Competition Commissioner that such combination would not be possible. [92] Recommending could potentially also be used by a platform with a market position with one service in one market to assist with leveraging market position with another service in an adjacent market. Using knowledge of user preferences, behaviours, interests, and so on derived from tracking on one service operating in one market, a new service in an adjacent market can potentially be more personalised, therefore more attractive to users, than might otherwise be the case. This could boost the adoption and therefore market position of that service, allowing the platform in question to more effectively move into new markets, increasing the sources from which data can be gathered and from which profit can be generated.

With the ability to promote or demote certain content, recommender systems also allow dominant platforms to position themselves as gatekeepers of political discussion, capable of facilitating or impeding the dissemination of information [93] (usually with few, if any, legal safeguards). Facebook's own research has twice demonstrated that small changes to the content shown to users by their recommender system can have an effect on election outcomes. [94] And Epstein and Robertson have demonstrated a 'search engine manipulation effect', whereby manipulating the rankings of political search results can not only be successfully masked but can shift the voting preferences of undecided voters by 20 per cent or more. [95] Tambini points out that, in the political arena, "Facebook, in particular, is emerging as a vertically integrated one-stop-shop for fundraising, recruitment, database building, segmentation, targeting, and message delivery", [96] and that the dynamics of electorates and elections are now more knowable than ever - to Facebook , a private company, and to anyone willing to pay, but not necessarily to other parties or to election regulators. Indeed, in the 2016 US Presidential election, various platform companies were involved in helping campaigns shape their message and plan and execute digital strategies. [97]

Zittrain warns of the potential for what he calls 'digital gerrymandering' - "the selective presentation of information by an intermediary to meet its agenda rather than to serve its users". [98] The willingness of platforms to use their position to suit their political agenda is not hypothetical; Google, for example, has used its website to protest the US Stop Online Piracy Act [99] and to promote messages related to net neutrality and LGBT rights. [100] Their ability to undertake similar campaigns is one reason why newspapers are subject to merger and competition rules to limit market concentration. [101]

PRIVATE ORDERING

Recommending gives rise to extensive and influential new forms of private ordering, with platforms setting their own rules for recommending, their own criteria for constructing online spaces, and their own mechanisms for dealing with some of the problems that arise. Online spaces are algorithmically constructed; as algorithms are inherently normative, whoever operates those algorithms has enormous influence over that process and thus over those spaces. This influence is amplified by monopolisation and the growing power of platforms. Continuing to allow platforms a large degree of self-regulation in relation to recommending allows them to establish their own standards, guidelines, and enforcement mechanisms. [102] As a result, terms of service, community standards, and the code underpinning the platforms itself effectively become the primary law on those platforms, [103] enforced by the platform's own review bodies. Indeed, Facebook has proposed to create its own 'supreme court' to adjudicate on issues arising on its platform. [104]

This gives platforms much influence over what speech is acceptable to be algorithmically disseminated across the internet - now a key forum for the public sphere - and how that acceptability is policed, with little transparency in how rules and mechanisms are established and maintained, often little accountability to users, and usually without oversight. The use of recommender systems by platforms to construct online space is primarily motivated by profit and duty to shareholders rather than concern for public good and responsibility to wider society. It is plainly undesirable for this influence to be placed in the hands of private corporations without minimum standards and with little oversight.

As a result, the influence over the public sphere generated by recommending should be reduced by developing a legal framework applying equally across all platforms to establish uniform minimum rules, standards, and accountability and enforcement mechanisms for content dissemination by recommending.

4. Current Law and Liability

In the EU, the principal law in this space is currently the E-Commerce Directive. [105] The Directive provides for certain liability protections in relation to user-generated content for providers of information society services(as defined in the Technical Standards and Regulations Directive: "any service normally provided for remuneration, at a distance, by electronic means and at the individual request of a recipient of services" [106])who are acting as intermediary service providers. [107] These protections apply to three specified activities: acting as a mere conduit (transmitting information between individuals - e.g. ISPs and messaging services) [108] ; caching (a technical activity undertaken in relation to acting as a conduit which assists with the efficiency of the service) [109] ; and hosting (storing information provided by a recipient of the service in question) [110]. In relation to acting as a mere conduit and caching, service providers are protected from liability for content (subject to several conditions). [111] For hosting, they are protected so long as they have no actual knowledge of illegal activity or information or where, upon obtaining such knowledge, they expeditiously remove or disable access to the content. [112] It should also be noted that while platforms often store recommended content on their own servers, this is not by itself sufficient to avail of the hosting protection. [113] For that to be available, the activity ("[consisting of] the storage of information provided by a recipient of the service" [114]) must be neutral with regard to the content in question and must passive and merely technical. [115]

To be acting as an intermediary service provider, a service provider must be neutrally providing a service by automatic, technical, and passive means. [116] The E-Commerce Directive, in Recital 42, states that the Directive's liability protections are limited to situations where the service provider has engaged in an activity of a "mere technical, automatic and passive nature, which implies ... neither knowledge of nor

control over the information which is transmitted or stored". [117] The CJEU has interpreted this to mean that service providers will not be intermediary service providers (and will thus not be afforded liability protection) where they (a) take an active role in relation to content that would (b) give them either knowledge of or control over that content. [118] Where service providers promote certain content, for example, they will not have taken a neutral role between uploader and viewer as an intermediary service provider but will have (a) taken an active role that (b) gives them knowledge of or control over that content and will therefore not be able to rely on the Directive's liability protections. [119] Following Recital 42 and the CJEU's reasoning, it is not sufficient to have (a) without (b), or vice-versa; in order for a service provider to be excluded from the Directive's liability protections, both must be made out. Ascertaining whether any protection applies must therefore involve determining whether this is indeed the case. [120]

Straightforwardly, no liability protections apply to recommending of content produced by a platform itself, regardless of the form of recommending involved. Such protections are available only in relation to information provided by a recipient of the service, [121] not content produced by the service provider. While closed recommending is therefore excluded from protection, curated recommending could potentially involve user-generated content, and open recommending almost certainly will. However, even if curated recommending does include user-generated content, this does not involve the platform acting as an intermediary service provider. In curated recommending, by its nature, they will have actively and deliberately selected the content in question and will therefore likely have both knowledge of and control over that content.

The situation is somewhat different in relation to open recommending of user-generated content, as service providers are unlikely, in most cases, to have knowledge of the content itself. Rather, they will have knowledge of metadata about the content - for example, information about which users have viewed or provided feedback or indications of its general 'popularity' (e.g. 'likes' or 'shares'). However, in open recommending, service providers do exercise control over content. Indeed, control over content is the very point of recommending; service providers exercise such control in order to show people what they want them to see in order to drive engagement, profit, and market position. The normative nature of recommender systems - the fact that they enable platforms to exercise control over content distribution in pursuit of their own goals - is the reason for their use.

In open recommending, service providers are not simply storing user content and displaying it neutrally in a merely technical or passive way. They control the criteria for recommending and thus they determine what is recommended in line with the outcomes that they desire. According to the CJEU, simply providing general information to users cannot itself be sufficient to deprive a service provider of the protection afforded to hosts. [122] But, in recommending content, service providers do not provide general information - they provide and promote specifically selected information, determined by the service provider (by way of their algorithmic processes) on the basis of predicted relevance, interest, and so on to the user or groups of users in question. They therefore do not take a neutral position between the uploader of the content and the potential viewers [123] ; rather, they are actively involved in selecting content for distribution and promotion according to their own criteria, in selecting the audience for that content according to their own determination, and, as a result, in shaping online spaces for their own purposes and benefit. In doing so, they take an active role of a kind to give them control over the content in question. Consequently, in relation to open recommending, platforms cannot be intermediary service providers and are operating beyond the limits of the liability protections provided by the E-Commerce Directive.

The effect of this analysis would be that recommending (whether open, curated, or closed) is not an activity that is covered by any of the E-Commerce Directive's liability protections. This argument has received some tentative support in the domestic courts. The Italian Supreme Court has gone so far as to

enumerate activities that would take service providers outside of the Directive's liability protections on grounds of enriching the use of content by undetermined users in a non-passive way. [124] These would include, among other things, filtering, selecting, organising, classifying, aggregating, or promoting content as part of the business activities of the service provider, as well as the adoption of behaviour assessment techniques intended to increase user retention. [125] The High Court of England and Wales has gone further and has suggested - albeit in obiter - that simply considering content for recommending might also take service providers outside of the Directive's protections even where that content is not subsequently actually recommended. [126]

Recommending is an activity engaged in by service providers - distinct from hosting, caching, or acting as a mere conduit - which developed in the years after the E-Commerce Directive was passed. However, this distinct activity of recommending is also not otherwise covered under any particular regulation beyond the general law. This means that recommending falls into a significant and consequential gap in the current legal regime. Following this analysis, service providers would have, under the E-Commerce Directive or any other current law, no special protection against liability for recommended illegal content or activity on their platforms. [127] This requires attention. Due to the risk of piecemeal or incoherent legal responses to this gap emerging from the various Member States or elsewhere, the development of a considered and comprehensive regulatory response is plainly desirable.

5. Responsible Recommending

The previous analysis identified a consequential gap in the current legal framework governing the liability protections of service providers. The straightforward response would be to provide for a substantively similar regime to that already established for hosting, whereby service providers are shielded from liability provided they expeditiously remove illegal recommended content once they become aware of it. However, recommending is a much less neutral, more involved, and more active form of service provider activity than simply hosting. While liability protection for recommending is perhaps desirable, this should come with some responsibilities beyond simply removing illegal content expeditiously upon acquiring knowledge of the illegality. These responsibilities would sit alongside and complement other applicable legal frameworks, such as around data protection or non-discrimination. Regulating recommending therefore provides an opportunity to establish a more inclusive framework of principles, requirements, and limitations within which the discrete activity of recommending can responsibly be undertaken.

In line with this, recent regulation passed by the EU focuses on recommending by platforms of goods and services offered by business users through the platform, [128] primarily aiming to provide for transparency and to improve fairness between market participants in that business environment. [129] Focusing on the "platform-to-business" [130] relationship, it may go some way towards limiting the ability of service providers to leverage their influence over recommender systems' content rankings to effect changes in markets which have come to rely on their platform. However, most of the issues identified in this paper do not relate to online markets for good and services. Further regulation is therefore required to address the deficiencies in the current legal regime around recommending, as well as to provide considered responses to the other issues arising from recommending discussed above.

Following the above analysis, three areas where recommending makes a significant contribution to systemic issues can be identified: hate speech and violent extremism; disinformation and conspiracy theories; and monopolisation and platform power. Legal responses to these problems are not straightforward given the various interests to be considered and the fact that platforms typically process large volumes of information. Any form of content moderation or restriction on recommendations is difficult to automate, culturally contextual, and potentially sensitive. But if it is accepted that there are

issues with recommending as currently undertaken then the law should consider the best (or, as it may be, least bad) approach to preventing or mitigating those issues to at least some meaningful extent.

As previously discussed, regulation focusing on the transmission or hosting of content itself brings freedom of expression concerns. However,the same risks do not necessarily arise from regulating the further dissemination of content by corporations. While the fundamental right of freedom of expression should be respected as far as possible, individuals do not have a fundamental right to have their speech disseminated or amplified by platforms in this way (nor would it be desirable to establish a right to be recommended akin to a right to a platform). By focusing on recommending, rather than on the transmission or hosting of content itself, regulation can largely sidestep these freedom of expression problems and focus on the use of technical systems by private corporations to pursue their own business goals. No obligation on the part of service providers to monitor, identify, remove, or prevent the upload of content would arise from these principles, and service providers would not face any liability for legal but potentially problematic content that is submitted by users. Individuals would remain free to post, search for, browse for, view, or share any content which is lawful. The onward algorithmic dissemination and amplification of such content by service providers would be the focus of regulation.

Seemingly obvious approaches to dealing with some of the issues discussed above may not have the desired effect. For example, it might be thought that requiring content diversity - as has been proposed by regulators in Germany [131] - would assist with addressing disinformation, extremism, and fragmentation. However, this could mean that disinformation and conspiracy theories would be actively promoted alongside factual information, as if they were all of equal value. Yet research suggests that exposure to disinformation increases the likelihood of it being perceived as accurate [132]; that in some cases showing people things they disagree with may actually increase polarisation [133]; and that, in any case, people tend to select things that support or confirm their existing views, which can further polarise. [134] As a result, requiring content diversity might actually be a counterproductive strategy for addressing some of these issues.

Taking into account the issues discussed above, it is possible to set out some general principles for responsible recommending that can inform a future legal framework (indeed, setting minimum standards is in and of itself a response to private ordering; some of these proposed principles adopt existing practices undertaken by one or more service providers into minimum standards applying generally). This is not an exhaustive list; these high-level principles do not address all of the issues identified previously, and there are likely others that are worth adopting. These principles would not represent radical interventions but considered responses that could go some way towards mitigating the systemic problems to which open recommending can contribute.

The first two principles for responsible recommending proposed below would establish a general regulatory framework for liability protection to be granted on a conditional basis for open recommending. The remaining four principles set out some more specific requirements for service providers operating within that framework to obtain or lose that protection. These principles do not establish a duty of care to individuals, nor a more general duty of care to provide a safe environment. [135] They do not require awkward analogies with publishers or public spaces, [136] and establish no general obligation to prevent 'harmful' content or activities by users. [137] Instead, they propose a requirement to recommend responsibly, including an obligation to work towards reducing and eliminating the recommending of potentially problematic content, with a fall-back to a platform-specific prohibition on open recommending if service providers fail to fulfil their requirements arising from these principles. The difficulty of defining various terms like 'disinformation', 'conspiracy theories', and 'violent extremism' means that law should not establish liability for recommending content itself. These principles do not therefore establish liability for legal but potentially problematic content instead; instead, they provide for responsibilities to work towards reducing the dissemination of potentially problematic

content through recommending, restrictions for systematic failures to meet those responsibilities, and potential liability for failure to obey those restrictions.

1. Open recommending must be lawful and service providers should be prohibited from doing it where they violate these principles or other applicable laws. The problems caused or exacerbated by recommender systems largely manifest with open recommending for engagement, revenue, and market position without due attention being paid to consequences and side effects in the shape of the issues discussed above. Seeking to address this, this principle reflects the idea that if service providers can't do open recommending responsibly then they shouldn't be permitted to do it at all.

Beyond the general principles for responsible recommending set out herein, other applicable legal frameworks must also be considered. Data protection, especially, is of fundamental importance in this context, given the extensive behavioural tracking and processing of personal data which underpins recommending. [138] Responsible recommending cannot be undertaken while flouting data protection law. [139] Nor can service providers ignore their responsibilities under equality and non-discrimination law [140] (or, indeed, any other applicable regime).

Should a service provider's systematic and repeated failures to fulfil any of these principles or to meet other applicable laws point to an inability or unwillingness to do so, the service provider should be prohibited from open recommending. This prohibition should be imposed until the service provider can demonstrate that they are in a position to adequately meet their responsibilities. Where they are prohibited from open recommending, service providers could face financial penalties if they continue to use open systems and should not be able to avail of the liability protections normally afforded to open recommending by the second principle. This would incentivise service providers to employ mechanisms to deal with these principles, in particular around recommending potentially problematic content, and would go some way towards preventing service providers who are unable to recommend responsibly from contributing to systemic problems, while providing a route to the prohibition being lifted should a service provider acquire that capability.

2. Service providers should have conditional liability protections for recommending illegal user-generated content and should lose liability protection for recommending while under a prohibition. As argued previously, in relation to recommending, service providers are not currently covered by the liability protections established in the E-Commerce Directive. These liability protections should be extended to cover user-generated or externally-sourced content that is recommended in open systems, provided illegal content is removed expeditiously upon otherwise obtaining actual knowledge or awareness of the illegality in much the same way as is required for hosting. However, liability protections should not be extended to cover curated or closed systems, as they involve content that has either been produced or deliberately selected (for example, through licensing) by the service provider. Liability protection should be lost where service providers systematically fail to recommend responsibly and, consequently, are prohibited from open recommending. No liability protection would therefore be available when undertaking a prohibited practice.

3. Service providers should have a responsibility to not recommend certain potentially problematic content. This is not a duty to prevent 'harm', nor is it an obligation to prevent certain kinds of content from being posted or hosted, and it is not liability arising from legal but potentially problematic content itself (whether hosted or recommended). This instead establishes a responsibility to not promote certain kinds of content (for example, white supremacism, health disinformation, anti-Semitic conspiracy theories, pro-suicide or self-harm content, content promoting eating disorders, and so on). The obvious challenge with this principle lies in the difficulty of defining terms in such a way as to include as much as possible of the content that should be included while including as little as possible of the content that shouldn't be. While every effort should be made to arrive at workable definitions, there are inevitably

going to be grey areas and the presumption should be in favour of not restricting the recommending of content where its classification is uncertain.

The E-Commerce Directive prohibits Member States from imposing a general obligation upon service providers to monitor content in the context of any of the three activities set out in Articles 12 to 14 (acting as a mere conduit, caching, and hosting). [141] But, as argued above, recommending is not an activity covered by Articles 12 to 14, and so would not come directly under this prohibition. That said, requirements that could in effect become such an obligation to monitor in relation to Articles 12 to 14 would also likely not be permitted. [142] However, under this principle, service providers would face no general obligation to identify illegal activity or remove illegal content. They would also not be liable for legal but potentially problematic content that they host or recommend, nor ordinarily for recommending illegal content where it is removed expeditiously upon subsequently obtaining actual knowledge or awareness of that illegality (as per the second principle). Instead, due to their active role in disseminating content for their own purposes, they would be responsible for avoiding recommending such content and, through systematic failures to meet that responsibility, would risk a prohibition on open recommending and the resulting loss of recommending liability protections. Any potential liability would therefore actually result from undertaking a prohibited practice - open recommending while under a prohibition - rather than from recommending certain kinds of potentially problematic but lawful content.

4. Service providers should be required to keep records and make information about recommendations available to help inform users and facilitate oversight. Research suggests that a majority of users may not be aware that the Facebook News Feed is algorithmically constructed and do not understand how it works. [143] For example, Eslami et al found that 62.5% of Facebook users studied did not know that Facebook's News Feed is not chronological, and that many of these users had incorrectly inferred things about their social relationships as a result of certain individuals not appearing in their feeds. [144] Similar 'folk theories' about News Feed, relying on incomplete or non-existent understandings of its algorithmic nature, have been observed elsewhere. [145] The general lack of transparency around recommending is also a problem in terms of oversight, making it difficult to determine what material has been recommended to whom.

Transparency may therefore seem like an obvious solution to inform and, ideally, empower users. But future law should be wary of falling into the 'transparency fallacy', [146] whereby transparency seems like a remedy but in fact merely gives individuals unhelpful information and fails to provide the anticipated benefits of empowerment and control. Merely providing information can be used to manipulate users into trusting a system to their detriment [147] ; indeed, Facebook's existing advert explanations are often incomplete, vague, and misleading. [148] Transparency cannot solve problems and is not a cure, [149] but it can be one potentially helpful element supporting a broader regime which itself can be more effective. In an online context, transparency is about more than just end-users, but also having information available for oversight agencies, regulators, civil society organisations, informed minorities, and so forth. [150] That is to say, while not a solution itself, transparency can be considered to be a general principle to facilitate other principles and oversight.

The constantly changing nature of social media and other online services makes it difficult to identify and track problems over time. Service providers should, at a minimum, be required to keep logs of recommended content (both for personalisation and for behavioural targeting) so that they can be reviewed by users and by oversight bodies (for example, a former Google engineer has created a publicly available tool that shows which videos YouTube has been recommending each day [151] - the platforms themselves would be in a position to provide far more accurate and detailed information). The nature of the records and how they are presented may vary to suit the audience; simple and summary information can be made generally available, while full and detailed information should be available to regulators and interested users. Service providers should also be required to provide information to users on where

content has come from and why it has been recommended (along these lines Facebook has incorporated a 'why am I seeing this' tool in News Feed [152]).

5. Recommending should be opt-in, users should be able to exercise a minimum level of control over recommending, and opting-out again should be easy. This is linked to transparency, as transparency can help users make informed choices, but again should not be regarded as a panacea and regulation should be careful to not overload users with options. Doing so would likely will result in something akin to the 'privacy paradox', whereby individuals profess to care about privacy but routinely fail to take steps to protect it (as a result of being overloaded with options [153] and perhaps now wary of privacy policies and controls [154]). That said, offering control is still a good idea, provided those who do not exercise it do not end up being treated less favourably than those who do. To that end, it would be desirable for recommending to be available to users only on an opt-in basis. Users who choose to receive recommendations should, at a minimum, be able to exclude certain content from recommendations, should be able to exclude certain sources of content from recommendations, should be able to exclude certain of their behaviours or interests from the process of determining what should be recommended to them, and should be able to easily and freely opt back out of recommendations entirely. Such controls might also facilitate interoperable and transferable control policy templates, [155] allowing certain settings to be defined by one user and easily adopted by others. This would allow, for instance, civil advocacy groups and so forth to define settings that users could easily adopt.

6. There should be specific restrictions on service providers' ability to use recommending to influence markets through recommending. Service providers should be explicitly prohibited from unduly recommending their own products and services ahead of those offered by others (it is worth noting that India has gone further and has prohibited e-commerce companies, including Amazon, from selling their own products on their websites). [156] As discussed previously, the European Commission has already fined Google for doing this; these developments would be brought into a future framework regulating recommending (similar to how GDPR brought the so-called 'right to be forgotten', established in the CJEU decision in Google Spain [157] , expressly into the data protection framework [158]).

Further, service providers (as data controllers) could be prohibited from using personal data obtained from acquisitions in their main business recommender systems without, at a minimum, explicit consent (as the German competition authority has ordered for Facebook). Likewise, service providers could be prohibited from using knowledge of user interests and preferences derived from their analysis of user behaviour in one service when they move into an adjacent market with a new service. These prohibitions would not only complement and refine the existing data protection principle of purpose limitation, [159] but would go some way towards addressing competition issues arising from the dominance of platforms and their use of personal data, particularly where leveraging their dominant position in one market to gain a competitive advantage in another.

6. Oversight and Compliance

The principles for responsible recommending proposed in this paper would require some body to assess evidence of success at fulfilment, with the power to impose prohibitions and any other enforcement mechanisms where necessary. This body, however constituted, would ideally be independent of government (so as to mitigate against the potential for political interference), and should perhaps itself be under judicial, legislative, or similarly appropriate supervision. Service providers and oversight bodies should consider things in aggregate, identifying patterns and trends in recommending rather than focusing on individual items of content.

Compliance should not be assessed on a zero-tolerance basis - oversight should acknowledge that this is difficult, that definitions are not straightforward, that no system is perfect, and that it is inevitable that some potentially problematic content will slip through. The regulatory focus would primarily be on ensuring that service providers act in good faith, are actually employing measures to try to get things right, and are in fact making progress towards more responsible recommending. To this end, the principles above take a 'carrot and stick' approach to incentivise ongoing progress towards compliance by service providers, rather than a purely sanctions-based one. The liability protections granted for responsible recommending combined with the risk of prohibition and the suspension of those protections for systematic failures should incentivise service providers to develop better (if perhaps still imperfect) practices. That said, oversight bodies must be prepared to actually impose and enforce the prohibition on recommending where service providers fail to meet their responsibilities - they cannot be thought to be optional.

Something along the lines of a co-regulatory approach may be useful here. Co-regulation emphasises a dialogue between parties working together to achieve the desired outcomes, [160] but typically lacks both a statutory basis and state-backed sanctions for non-compliance. A solution may involve adopting a collaborative and dialogue-based approach derived from co-regulation but underpinned by legislation providing for general definitions, principles, and minimum standards, qualified liability protections for open recommending, and the power to prohibit a service provider from undertaking open recommending (and thus suspend liability protections). This would involve establishing guidelines, statutory codes of practice, and so on within this general legislative framework and then working with service providers, expert bodies, and other stakeholders to identify specific forms of potentially problematic content and to help service providers work towards meeting the principles for responsible recommending rather than simply hammering them with requirements and demanding immediate compliance. This could assist in making the connection between general definitions and principles and specific kinds of content to not be recommended, could help oversight and enforcement avoid being tied to overly precise and rigid definitions, and would allow for a more responsive regulatory regime.

Beyond that, this paper will not propose oversight or enforcement mechanisms, whether involving full (state) regulatory or co-regulatory approaches. Enforcement is a complicated issue involving many interested parties. Much work is required to figure out what, if any, existing oversight regimes might cover this area in different ways, where there are overlaps and gaps, which forms of review and enforcement would be most effective, how to deal with cross-border issues and extraterritoriality, and how to ensure adequate funding and expertise. In short, oversight and enforcement are difficult areas and simply proposing one mechanism or framework without fully considering all the issues involved would be of no benefit. Suffice it to say that this paper proposes some principles to be considered and implemented in future regulation, whatever shape the oversight and enforcement regime of that regulatory framework takes. However, while effective enforcement is necessary, it is worth nothing that it is likely not possible for one country acting on its own to do this effectively. Although oversight and enforcement may operate at a national level, the global nature of online platforms means that the establishment of legal standards for recommending requires coordinated action at supranational or international level if it is to have a realistic chance of having any significant effect.

COMPLIANCE AND IMPLEMENTATION

In some cases, the principles for responsible recommending reflect the patchwork of tools, mechanisms, and restrictions already implemented by various platforms. As a result, implementation of much of what this paper proposes should be more than feasible, and requirements such as logging and record-keeping to facilitate oversight would also be useful for service providers as they seek to recommend responsibly. That said, it may be that compliance with the responsibility to not recommend certain content is difficult

enough to achieve with open recommending that service providers move away from that. Even so, personalisation could and would still exist - recommending is not the only form it takes. Platforms that use curated or closed recommending to personalise services would essentially be able to carry on as before. Even under a prohibition on open recommending, service providers would still be able to show users a chronological or otherwise unfiltered feed of content produced by accounts they follow (as is currently an option on Facebook and Twitter, for example). And, of course, users would still be able to search for, browse for, and share content (either on the platforms themselves or elsewhere).

Regulating for responsible recommending may even drive the development of new practices by service providers. For example, a new form of semi-curated recommending could emerge. This could involve the platform choosing certain sources of content which would automatically feed into the pool of content from which the recommender system selects (for example, user accounts meeting certain criteria or that are well-established but have no history of producing potentially problematic content). However, in such a system, content produced by other accounts would not automatically be available for recommending. While not foolproof, this approach to recommending would be somewhere between open (in that service providers could benefit from liability protections as they wouldn't be specifically selecting or reviewing content itself, as in curated recommending, only sources) and curated (in that service providers wouldn't be exposed to the same level of risk inherent in recommending all user-generated content by default). While they could still face prohibitions for systematic failures, service providers using such a hybrid form would likely find it easier to responsibly recommend. Beyond law as a driver itself, negative press and social media coverage about which content is being recommended to whom has on several occasions in recent years led advertisers to stop spending money when paired with potentially problematic content. [161] Responsible recommending could help prevent this, providing a further business motivation for compliance.

This may come at a financial cost to platforms. But their profit-driven interest in recommending irresponsibly cannot outweigh society's interest in mitigating the problems to which it contributes. These businesses would not be the first to face potentially costly but societally necessary regulation; over decades, many others have had their practices and profits reined in for the greater good. [162] If service providers process too much information to be able to meet their responsibilities, then they should process less. If their service moves too fast, then it should slow down. If they are unable or unwilling to do what is necessary to meet their responsibilities, then they should face the prohibition described above. The inevitable costs and difficulties associated with responsible recommending should become understood by service providers and others as the cost of doing business.

7. Limitations and context

The problems arising from the internet are the result of a complex combination of societal and technical factors. Irresponsible recommending by platforms is just one of those factors, and a truly holistic response is needed to adequately address these issues. Regulating recommending is not in any way sufficient on its own and is not going to solve all of the problems with the internet. But responsible recommending could make a positive contribution towards a holistic set of solutions that could address many issues. Responsible recommending therefore needs to be understood as only one part of the much broader move that is needed to address problems and 'online harms' (however conceived of) and reduce the influence of platforms. This may involve interventions from competition law, [163] considering decentralisation and peer-to-peer communications, alternative approaches to data governance, [164] stronger protections for users, [165] better enforcement of data protection requirements, new technical means for control and audit, [166] and so on. Beyond this, as part of a more

holistic approach, education is important - the need to teach critical thinking and media literacy is key to combatting disinformation and conspiracy theories [167], for example.

8. Conclusions and further research

Many of the systemic issues arising on the internet stem from the irresponsible recommending of content by platforms in pursuit of engagement, profit, and market position. Many proposed regulatory responses to 'online harms' focused on content miss the fact that when it comes to more systemic problems the real issue is arguably recommending, as this is drives content dissemination. A focus on recommending would assist in the development of a legal framework which respects the right to freedom of expression by not restricting content production or hosting, while acknowledging that, in recommending content, platforms are often operating beyond the limits of where liability protections are currently provided to intermediary service providers.

Addressing recommending as an activity through principles-based regulation allows for the development of technology-neutral and platform-neutral legal responses. Compliance in terms of moving towards responsible recommending should be recognised as difficult, particularly in relation to the third principle (restricting the recommending of potentially problematic content), and oversight should adopt a collaborative, dialogue-based approach towards enforcement. A key goal of the high-level principles proposed in this paper is to provide incentives for service providers to work towards compliance, with liability protections for responsible recommending, on one hand, and the risk of prohibitions on open recommending and losing that protection, on the other.

[1] Compliant and Accountable Systems Group, Department of Computer Science and Technology, University of Cambridge (jennifer.cobbe@cl.cam.ac.uk / jatinder.singh@cl.cam.ac.uk). We acknowledge the financial support of the University of Cambridge (through the Cambridge Trust & Technology Initiative), the UK Engineering and Physical Sciences Research Council (EPSRC) [EP/P024394/1, EP/R033501/1], and Microsoft (through the Microsoft Cloud Computing Research Centre).

[2] Nick Srnicek (2016) Platform Capitalism , Polity Press

[3] Patrick Barwise and Leo Watkins (2018) 'The evolution of digital dominance: how and why we got to GAFA', In Martin Moore and Damian Tambini (eds.) Digital Dominance: The Power of Google, Amazon, Facebook, and Apple , Oxford University Press. Available at http://lbsresearch.london.edu/914/ [accessed 14/04/2019]

[4] Shoshana Zuboff (2019) The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power , Profile Books

[5] Directive 98/34/EC of the European Parliament and of the Council of 22 June 1998 laying down a procedure for the provision of information in the field of technical standards and regulations (Official Journal L 204 , 21/07/1998 P. 0037 - 0048) ('Technical Standards and Regulations Directive'), Art.1 (as amended by Directive 98/48/EC of the European Parliament and of the Council of 20 July 1998 amending Directive 98/34/EC laying down a procedure for the provision of information in the field of technical standards and regulations (Official Journal L 217 , 05/08/1998 P. 0018 - 0026)); Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000 on certain legal aspects of information society services, in particular electronic commerce, in the Internal Market (Official Journal L 178 , 17/07/2000 P. 0001 - 0016) ('E-Commerce Directive'), Recital 18

[6] E-Commerce Directive, Article 2(b); this paper uses the terms 'platform' and 'service provider' interchangeably

[7] Department for Digital, Culture, Media & Sport (2019) Online Harms White Paper , CP 57. Available at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/793360/Online_Harms_White_Paper.pdf [accessed 14/04/2019]

[8] Department for Digital, Culture, Media & Sport (2019), pp.30-32

[9] Department for Digital, Culture, Media & Sport (2019); House of Lords Select Committee on Communications (2019) Regulating in a Digital World , HL Paper 299. Available at https://publications.parliament.uk/pa/ld201719/ldselect/ldcomuni/299/299.pdf [accessed 14/04/2019]; Lorna Woods, William Perrin, and Maeve Walsh (2019) Internet Harm Reduction , Carnegie Trust. Available at https://www.carnegieuktrust.org.uk/publications/internet-harm-reduction [accessed 14/04/2019]; COM (2016) 593: Proposal for a Directive of the European Parliament and of the Council on copyright in the Digital Single Market

[10] See, e,g., TJ McIntyre (2018) 'Internet Censorship in the United Kingdom: National Schemes and European Norms' In Lilian Edwards (ed.) Law, Policy and the Internet , Hart Publishing. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3182549 [accessed 14/04/2019]; Ben Wagner (2016) Global Free Expression - Governing the Boundaries of Internet Content , Springer. Available at https://www.springer.com/gb/book/9783319335117 [accessed 14/04/2019]

[11] Reuben Binns, Michael Veale, Max Van Kleek, Nigel Shadbolt (2017) 'Like trainer, like bot? Inheritance of bias in algorithmic content moderation', 9th International Conference on Social Informatics (SocInfo 2017) . Available at https://arxiv.org/abs/1707.01477 [accessed 27/07/2019]; Brian Merchant (2019) 'How a Horrific Murder Exposes the Great Failure of Facebook's AI Moderation', Gizmodo . Available at https://gizmodo.com/the-great-failure-of-facebook-s-ai-content-moderation-s-1836500403 [accessed 27/07/2019]

[12] Sarah T Roberts (2014) 'Behind the screen: the hidden digital labor of commercial content moderation', PhD Thesis . Available at https://www.ideals.illinois.edu/handle/2142/50401 [accessed 14/04/2019]; Sarah T Roberts (2016) 'Digital Refuse: Canadian Garbage, Commercial Content Moderation and the Global Circulation of Social Media's Waste', Wi: Journal of Mobile Media . Available at https://ir.lib.uwo.ca/commpub/14/ [accessed 14/04/2019]

[13] Barwise and Watkins (2018); Ulrich Dolata (2017) 'Apple, Amazon, Google, Facebook, Microsoft: Market concentration - competition - innovation strategies', Stuttgarter Beiträge zur Organisations- und Innovationsforschung, SOI Discussion Paper, No. 2017-01 . Available at https://ideas.repec.org/p/zbw/stusoi/201701.html [accessed 14/04/2019]

[14] Dolata (2017)

[15] Zuboff (2019)

[16] Council of Europe (2019) Declaration by the Committee of Ministers on the manipulative capabilities of algorithmic processes , art.8. Available at https://search.coe.int/cm/pages/result_details.aspx?ObjectId=090000168092dd4b [accessed 14/04/2019]

[17] House of Lords Select Committee on Communications (2019) para 193

[18] There may be various other factors, such as the specifics of the database management infrastructure, that could impact search results (for example, their ordering). This paper does not explore such technical detail, instead noting that these factors may be relevant if they form part of the platform's approach to influence query results.

[19] Engin Bozdag (2013) 'Bias in algorithmic filtering and personalisation', Ethics in Information Technology , 15, pp.209-227. Available at https://link.springer.com/article/10.1007/s10676-013-9321-6 [accessed 14/04/2019]; Tarleton Gillespie (2014) 'The Relevance of Algorithms', In Tarleton Gillespie, Pablo J Boczkowski, and Kirsten A Foot (eds.) Media Technologies: Essays on Communication, Materiality, and Society , MIT Press; Robin Hill (2016) 'What An Algorithm Is', Philosophy and Technology , 29(35). Available at https://link.springer.com/article/10.1007/s13347-014-0184-5 [accessed 14/04/2019]; David Beer (2017) 'The Social Power of Algorithms', Information, Communication & Society , 20(1). Available at http://eprints.whiterose.ac.uk/104026/1/Algorithms_editorial_final.pdf [accessed 14/04/2019]; Natascha Just and Michael Latzer (2017) 'Governance by algorithms: reality construction by algorithmic selection on the Internet', Media, Culture & Society, 39(2), pp.238-258. Available at https://journals.sagepub.com/doi/abs/10.1177/0163443716643157?journalCode=mcsa [accessed 14/04/2019]

[20] Beer (2017) p.4

[21] Nicholas Diakopoulos (2015) 'Algorithmic Accountability', Digital Journalism , 3(3), p.400. Available at https://www.tandfonline.com/doi/abs/10.1080/21670811.2014.976411 [accessed 14/04/2019]; Gillespie, 2014, 167; for a more complete analysis of what an algorithm is, see Hill, 2016, p.47

[22] Rob Kitchin (2017) 'Thinking critically about and researching algorithms', Information, Communication & Society , 20(1), pp.14-29. Available at https://www.tandfonline.com/doi/abs/10.1080/1369118X.2016.1154087?journalCode=rics20 [accessed 14/04/2019]; R Stuart Geiger (2014) 'Bots, bespoke, code and the materiality of software platforms', Information, Communication & Society , 17(3), pp.342-356. Available at https://www.tandfonline.com/doi/abs/10.1080/1369118X.2013.873069?journalCode=rics20 [accessed 14/04/2019]

[23] Shoshana Zuboff (2015) 'Big other: surveillance capitalism and the prospects of an information civilization', Journal of Information Technology , 30, pp.75-89

[24] Zuboff (2015); see also Mark Andrejevic (2011) 'Surveillance and Alienation in the Online Economy', Surveillance & Society , 8(3), pp.270-287; Christian Fuchs (2011) 'A Contribution to the Critique of the Political Economy of Google', Fast Capitalism , 8(1); Karl Palmås (2011) 'Predicting What You'll Do Tomorrow: Panspectric Surveillance and the Contemporary Corporation', Surveillance & Society , 8(3), pp.338-354

[25] Wolfie Christl (2017) 'Corporate Surveillance in Everyday Life', Cracked Labs . Available at https://crackedlabs.org/en/corporate-surveillance [accessed 14/04/2019]; Sarah Myers West (2017) 'Data Capitalism: Redefining the Logics of Surveillance and Privacy', Business & Society . Available at https://journals.sagepub.com/doi/full/10.1177/0007650317718185 [accessed 14/04/2019]; Wolfgang

Kerber (2016) 'Digital Markets, data, and privacy: competition law, consumer law and data protection', Journal of Intellectual Property Law & Practice , 11(11). Available at https://academic.oup.com/jiplp/article/11/11/856/2335247 [accessed 14/04/2019]

[26] As of April 2019 (Alexa Internet, Alexa Traffic Rank: Top 500 Sites . Available at https://www.alexa.com/topsites [accessed 14/04/2019]). Website ranking is crude and imperfect for a number of reasons, but provides a rough guide to which sites are most popular. Note that this ranking does not include traffic through apps.

[27] Jesus Bobadilla, Fernando Ortega, Antonio Hernando, and Abraham Gutiérrez (2013) 'Recommender systems survey', Knowledge-Based Systems , 46, pp.109-132. Available at https://www.sciencedirect.com/science/article/abs/pii/S0950705113001044 [accessed 14/04/2019]

[28] For a legally-accessible discussion of machine learning, see David Lehr and Paul Ohm (2017) 'Playing with the Data: What Legal Scholars Should Learn About Machine Learning' 51 U.C. Davis Law Review

[29] For example, Netflix (Carlos A Gomez-Uribe and Neil Hunt (2015) 'The Netflix Recommender System: Algorithms, Business Value, and Innovation', ACM Transactions on Management Information Systems , 6(4). Available at https://dl.acm.org/citation.cfm?id=2843948 [accessed 14/04/2019])

[30] In practice, it is common for open platforms to also include and recommend content they themselves have produced, commissioned, or acquired.

[31] Which may be covered by other legal redress mechanisms, such as in relation to IP infringement, defamation, hate crime offences, and so on

[32] Barwise and Watkins (2018); Dolata (2017)

[33] Nick Seaver (2018) 'Captivating algorithms: recommender systems as traps', Journal of Material Culture

[34] YouTube Creator Blog (2012) YouTube Now: Why We Focus on Watch Time . Available at https://youtube-creators.googleblog.com/2012/08/youtube-now-why-we-focus-on-watch-time.html [accessed 14/04/2019]

[35] Gomez-Uribe and Hunt (2015) p.7

[36] Qingyun Wu, Hongning Wang, Liangjie Hong, and Yue Shi (2017) 'Returning is Believing: Optimizing Long-term User Engagement in Recommender Systems', CIKM '17 Proceedings of the 2017 ACM on Conference on Information and Knowledge Management . Available at https://dl.acm.org/citation.cfm?id=3133025 [accessed 14/04/2019]

[37] Gomez-Uribe and Hunt (2015) p.5

[38] Ashley Rodriguez (2018) 'YouTube's recommendations drive 70% of what we watch', Quartz . Available at https://qz.com/1178125/youtubes-recommendations-drive-70-of-what-we-watch [accessed 14/04/2019]

[39] Barwise and Watkins, 2018; Dolata, 2017; Zeynep Tufekci (2016), 'As the Pirates Become CEOs: The Closing of the Open Internet', Dædalus, the Journal of the American Academy of Arts & Sciences , 145(1),

p.74. Available at https://www.mitpressjournals.org/doi/abs/10.1162/DAED_a_00366?journalCode=daed [accessed 14/04/2019]

[40] Sylvie Delacroix (2019) 'Beware of 'Algorithmic regulation''. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3327191 [accessed 14/04/2019]

[41] Just and Latzer (2017)

[42] Karen Yeung (2017a) 'Algorithmic Regulation: A Critical Interrogation', Regulation & Governance . Available at https://onlinelibrary.wiley.com/doi/full/10.1111/rego.12158 [accessed 14/04/2019]; partly in response to Yeung, Delacroix emphasises the difference between regulation and regulatory power (Delacroix, 2019)

[43] Jon Danaher (2016) 'The Threat of Algocracy: Reality, Resistance and Accommodation', Philosophy and Technology , 29, pp.245-268. Available at https://link.springer.com/article/10.1007/s13347-015-0211-1 [accessed 14/04/2019]

[44] Antoinette Rouvroy and Thomas Berns (2013) 'Algorithmic Governmentality and Prospects of Emancipation', Réseaux , 1(177). Available at https://www.cairn-int.info/article-E_RES_177_0163--algorithmic-governmentality-and-prospect.htm [accessed 11/05/2019]

[45] Just and Latzer (2017)

[46] Daniel Susser, Beate Roessler, Helen Nissenbaum (2019) 'Online Manipulation: Hidden Influences in a Digital World', p.2. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3306006 [accessed 14/04/2019]

[47] Gillespie (2014) p.167

[48] Tufekci (2016) pp.71-72; see also James G Webster (2010) 'User Information Regimes: How Social Media Shape Patterns of Consumption', Northwestern University Law Review , 104(2)

[49] Zeynep Tufekci (2015) 'Algorithmic Harms Beyond Facebook and Google: Emergent Challenges of Computational Agency', Colorado Technology Law Journal , 13, pp.207-208. Available at https://ctlj.colorado.edu/wp-content/uploads/2015/08/Tufekci-final.pdf [accessed 14/04/2019]

[50] Taina Bucher (2012) 'Want to be on top? Algorithmic power and the threat of invisibility on Facebook', New Media & Society , 14(7), pp.1164-1180. Available at https://journals.sagepub.com/doi/abs/10.1177/1461444812440159 [accessed 14/04/2019]; Taina Bucher (2017) 'The algorithmic imaginary: exploring the ordinary affects of Facebook algorithms', Information, Communication & Society , 20(1), pp.30-44. Available at https://www.tandfonline.com/doi/abs/10.1080/1369118X.2016.1154086 [accessed 14/04/2019]

[51] Adam D I Kramer, Jamie E Guillory, and Jeffrey T Hancock (2014) 'Experimental evidence of massive-scale emotional contagion through social networks', PNAS , 111(24). Available at https://www.pnas.org/content/111/24/8788 [accessed 14/04/2019]

[52] Karen Yeung (2017b) ''Hypernudge': Big Data as a mode of regulation by design', Information, Communication & Society , 20(1), pp.118-136. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2807574 [accessed 14/4/2019]

[53] Deliberately false information intended to mislead

[54] Julio Reis, Fabrício Benevenuto, Pedro O S Vaz de Melo, Raquel Prates, Haewoon Kwak, and Jisun An (2015) 'Breaking the News: First Impressions Matter on Online News', ICWSM 2015 . Available at https://arxiv.org/abs/1503.07921 [accessed 14/04/2019]; Soroush Voshougi, Deb Roy, and Sinan Aral (2018) 'The spread of true and false news online', Science , 359(6380), pp.1146-1151. Available at https://science.sciencemag.org/content/359/6380/1146.full [accessed 29/07/2019]; Paul Lewis (2018) ''Fiction is outperforming reality': how YouTube's algorithm distorts truth', The Guardian . Available at https://www.theguardian.com/technology/2018/feb/02/how-youtubes-algorithm-distorts-truth [accessed 14/04/2019]; Zeynep Tufekci (2018) 'YouTube, the Great Radicalizer', The New York Times . Available at https://www.nytimes.com/2018/03/10/opinion/sunday/youtube-politics-radical.html [accessed 14/04/2019]; Bryan Gardiner (2015) 'You'll Be Outraged At How Easy It Was To Get You To Click On This Headline', Wired . Available at https://www.wired.com/2015/12/psychology-of-clickbait [accessed 14/04/2019]

[55] Derek O'Callaghan, Derek Greene, Maura Conway, Joe Carthy, Pádraig Cunningham (2014) 'Down the (White) Rabbit Hole: The Extreme Right and Online Recommender Systems', Social Science Computer Review , 33(4), pp.459-478. Available at https://journals.sagepub.com/doi/abs/10.1177/0894439314555329?journalCode=ssce [accessed 14/04/2019]; Adrienne Massanari (2017) '#Gamergate and The Fappening: How Reddit is algorithm, governance, and culture support toxic technocultures', new media & society , 19(3), pp.329-346. Available at https://journals.sagepub.com/doi/abs/10.1177/1461444815608807 [acessed 14/04/2019]; Jonas Kaiser (2018) 'How YouTube helps to unite the Right', Alexander von Humboldt Institute for Internet and Society - Digital Society Blog . Available at https://www.hiig.de/en/how-youtube-helps-to-unite-the-right [accessed 14/04/2019]; Zoë Beery (2019) 'How YouTube reactionaries are breaking the news media', Columbia Journalism Review . Available at https://www.cjr.org/analysis/youtube-breaking-news.php [accessed 14/04/2019]; Lewis (2018); Kelly Weill (2018) 'How YouTube Built a Radicalization Machine for the Far-Right', The Daily Beast . Available at https://www.thedailybeast.com/how-youtube-pulled-these-men-down-a-vortex-of-far-right-hate [accessed 14/04/2019]; Caroline O'Donovan, Charlie Warzel, Logan McDonald, Brian Clifton, and Max Woolf (2019) 'We Followed YouTube's Recommendation Algorithm Down The Rabbit Hole', Buzzfeed News . Available at https://www.buzzfeednews.com/article/carolineodonovan/down-youtubes-recommendation-rabbithole [accessed 14/04/2019]

[56] Keegan Hankes and Alex Amend (2018) 'The Alt-Right is Killing People', Southern Poverty Law Centre . Available at https://www.splcenter.org/20180205/alt-right-killing-people [accessed 14/04/2019]

[57] Jessie Daniels (2018) 'The Algorithmic Rise of the "Alt-Right"', Contexts , 17(1). Available at https://journals.sagepub.com/doi/10.1177/1536504218766547 [accessed 14/04/2019]

[58] Massanari (2017)

[59] Kyle Wagner (2014) 'The Future of the Culture Wars Is Here, And It is Gamergate', Deadspin . Available at https://deadspin.com/the-future-of-the-culture-wars-is-here-and-its-gamerga-1646145844 [accessed 14/04/2019]

[60] George Hawley (2017) Making Sense of the Alt-Right , Columbia University Press, p.47; Kristin MS Bezio (2018) 'Crtl-Alt-Del: GamerGate as a precursor to the rise of the alt-right', Leadership , 14(5). Available at https://journals.sagepub.com/doi/abs/10.1177/1742715018793744 [accessed 14/04/2019]

[61] Casey Newton (2019) 'What the Covington Catholic debacle tells us about the internet today', The Verge . Available at https://www.theverge.com/2019/1/23/18193743/covington-catholic-social-media-internet-lessons [accessed 14/04/2019]

[62] Samantha Bradshaw and Phillip N Howard (2018) 'Why Does Junk News Spread So Quickly Across Social Media? Algorithms, Advertising and Exposure in Public Life', Oxford Internet Institute / Knight Foundation . Available at https://comprop.oii.ox.ac.uk/research/working-papers/why-does-junk-news-spread-so-quickly-across-social-media/ [accessed 14/04/2019]; John C Paolillo (2018) 'The Flat Earth Phenomenon on YouTube', First Monday , 23(12). Available at https://firstmonday.org/ojs/index.php/fm/article/view/8251/7693 [accessed 14/04/2019]; Julia Carrie Wong (2019) 'How Facebook and YouTube help spread anti-vaxxer propaganda', The Guardian . Available at https://www.theguardian.com/media/2019/feb/01/facebook-youtube-anti-vaccination-misinformation-social-media [accessed 14/04/2019] ; Matt Reynolds (2019) 'Think Facebook has an anti-vaxxer problem? You should see Amazon', Wired . Available at https://www.wired.co.uk/article/facebook-anti-vaccine-disinformation [accessed 14/04/2019]; Renee DiResta (2019) 'How Amazon's Algorithms Curated a Dystopian Bookstore', Wired. Available at https://www.wired.com/story/amazon-and-the-spread-of-health-misinformation [accessed 14/04/2019]

[63] Ysabel Gerrard and Tarleton Gillespie (2018) 'When Algorithms Think You Want to Die', Wired . Available at https://www.wired.com/story/when-algorithms-think-you-want-to-die [accessed 14/04/2019]

[64] Josephine B Schmitt, Diana Rieger, Olivia Rutkowski, and Julian Ernst (2018), 'Counter-messages as Prevention or Promotion of Extremism?! The Potential Role of YouTube', Journal of Communications , 68. Available at https://academic.oup.com/joc/article/68/4/780/5042003 [accessed 14/04/2019]

[65] Brandy Zadrozny (2019) 'Drowned out by the algorithm: Vaccination advocates struggle to be heard online', NBC News . Available at https://www.nbcnews.com/tech/tech-news/drowned-out-algorithm-pro-vaccination-advocates-struggle-be-heard-online-n976321 [accessed 14/04/2019]

[66] Eli Pariser (2012) The Filter Bubble: What The Internet Is Hiding From You , Penguin; see also Frederik Zuiderveen Borgesius, Damian Trilling, Judith Moeller, Balázs Bodó, Clas H. de Vreese, and Natalie Helberger (2015) 'Should We Worry About Filter Bubbles?', Internet Policy Review , 5(1a), pp.3-5 . Available at https://policyreview.info/articles/analysis/should-we-worry-about-filter-bubbles [accessed 14/04/2019]

[67] David Robert Grimes (2017) 'Echo Chambers are dangerous - we must try to break free of our online bubbles', The Guardian . Available at https://www.theguardian.com/science/blog/2017/dec/04/echo-chambers-are-dangerous-we-must-try-to-break-free-of-our-online-bubbles [accessed 14/04/2019]

[68] Lee De-Wit, Cameron Brick, Sander van der Linden (2019) 'Are Social Media Driving Political Polarization?', Greater Good . Available

at https://greatergood.berkeley.edu/article/item/is_social_media_driving_political_polarization [accessed 14/04/2019]

[69] Richard Fletcher and Rasmus Kleis Nielsen (2017) 'Are News Audiences Increasingly Fragmented? A Cross-National Comparative Analysis of Cross-Platform News Audience Fragmentation and Duplication' , Journal of Communication , 67(4). Available at https://onlinelibrary.wiley.com/doi/abs/10.1111/jcom.12315 [accessed 14/04/2019]; see also Borgesius et al (2015); Judith Möller, Damian Trilling, Natali Helberger, and Bram van Es (2018) 'Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity', Information, Communication & Society , 21(7), pp.959-977. Available at https://www.tandfonline.com/doi/full/10.1080/1369118X.2018.1444076 [accessed 14/04/2019]; Mario Haim, Andreas Graefe, and Hans-Bernd Brosius (2018) 'Burst of the Filter Bubble?', Digital Journalism, 6(3), pp.330-343. Available at https://www.tandfonline.com/doi/abs/10.1080/21670811.2017.1338145 [accessed 14/04/2019]

[70] Jianshu Weng, Ee Peng Lim, Jing Jiang, and Qi He (2010) 'Twitterrank: Finding topic-sensitive influential Twitterers', Proceedings of the Third ACM International Conference on Web Search & Data Mining: February 3-6, 2010, New York , pp. 261-270. Available at https://dl.acm.org/citation.cfm?id=1718520 [accessed 14/4/2019]; M D Conover, J Ratkiewicz, M Francisco, B Goncalves, A Flammini, and F Menczer (2011) 'Political Polarization on Twitter', Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media . Available at https://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/viewFile/2847/3275 [accessed 14/04/2019]; Antoine Boutet, Hyoungshick Kim, and Eiko Yoneki (2012) 'What's in Twitter: I know what parties are popular and who you are supporting now!', Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining , pp.132-129. Available at https://ieeexplore.ieee.org/document/6425772 [accessed 14/04/2019]; Robert Faris, Hal Roberts, Bruce Etling, Nikki Bourassa, Ethan Zuckerman, and Yochai Benkler (2017) 'Partizanship, Propaganda, and Disinformation: Online Media and the 2016 U.S. Presidential Election', Berkman Klein Center for Internet & Society Research Publication No. 2017-6 , p.71. Available at https://cyber.harvard.edu/publications/2017/08/mediacloud [accessed 14/04/2019]; Elanor Colleoni, Alessandro Rozza, and Adam Arvidsson (2014) 'Echo Chamber or Public Sphere? Predicting Political Orientation and Measuring Political Homophily in Twitter Using Big Data', Journal of Communication , 64, pp.317-332. Available at https://onlinelibrary.wiley.com/doi/10.1111/jcom.12084 [accessed 14/04/2019]; Eytan Bakshy, Solomon Messing, and Lada A Adamic (2015) 'Exposure to ideologically diverse news and opinion on Facebook', Science , 348(6239). Available at https://science.sciencemag.org/content/348/6239/1130 [accessed 14/04/2019]; Seth Flaxman, Shared Goel, Justin M Rao (2016) 'Filter Bubbles, Echo Chambers, and Online News Consumption', Public Opinion Quarterly , 80(S1). Available at https://academic.oup.com/poq/article/80/S1/298/2223402 [accessed 14/04/2019]

[71] Natalie Stroud (2010) 'Polarization and Partisan Selective Exposure', Journal of Communication , 60, pp.556-576. Available at https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-2466.2010.01497.x [accessed 14/04/2019]

[72] Ray Jiang, Silvia Chiappa, Tor Lattimore, András György, and Pushmeet Kohli (2019) 'Degenerate Feedback Loops in Recommender Systems', Proceedings of AAAI/ACM Conference on AI, Ethics, and Society, Honolulu, HI, USA, January 27-28, 2019 (AIES '19) . Available at https://arxiv.org/abs/1902.10730 [accessed 14/04/2019]; see also Nicola Perra and Luis E C Rocha (2018) 'Modelling Opinion Dynamics in the Age of Algorithmic Personalisation', arXiv Preprints , arXiv:1811.03341. Available at https://arxiv.org/abs/1811.03341 [accessed 14/04/2019]

[73] Borgesius et al (2016) p.8

[74] Mark Leiser (2016) 'AstroTurfing, 'CyberTurfing' and other online persuasion campaigns', European Journal of Law and Technology , 7(1). Available at http://ejlt.org/article/view/501 [accessed 14/04/2019]

[75] Leiser (2016)

[76] Leiser (2016)

[77] Leiser (2016); for a review of the literature, see Rose Marie Santini, Larissa Agostini, Carlos Eduardo Barros, Danilo Carvalho, Rafael Centeno de Rezende, Debora G Salles, Kenzo Seto, Camyla Terra, and Giulia Tuccy (2018) 'Software Power as Soft Power: A literature review on computational propaganda and political process', PArtecipazione e COnflitto , 11(2). Available at http://siba-ese.unisalento.it/index.php/paco/article/view/19546 [accessed 14/04/2019]; see also Samantha Bradshaw and Phillip N Howard (2017) 'Troops, Trolls and Troublemakers: A Global Inventory of Organized Social Media Manipulation', Computational Propaganda Research Project Working paper no 2017.12 , Oxford Internet Institute. Available at https://comprop.oii.ox.ac.uk/research/troops-trolls-and-trouble-makers-a-global-inventory-of-organized-social-media-manipulation/ [accessed 14/04/2019]; Bence Kollyani, Philip N Howard, and Samuel C Woolley (2016) 'Bots and Automation over Twitter during the U.S. Election', Data Memo 2016.4. Oxford, UK: Project on Computational Propaganda . Available at https://comprop.oii.ox.ac.uk/research/working-papers/bots-and-automation-over-twitter-during-the-u-s-election/ [accessed 14/04/2019]; Emilio Ferrara (2017) 'Disinformation and Social Bot Operations in the Run Up to the 2017 French Presidential Election', First Monday . Available at https://firstmonday.org/ojs/index.php/fm/article/view/8005/6516 [accessed 14/04/2019]; Alessandro Bessi and Emilio Ferrara (2016) 'Social bots distort the 2016 U.S. Presidential election online discussion', First Monday , 21(11). Available at https://firstmonday.org/article/view/7090/5653 [accessed 14/04/2019]; Muhammad Nihal Hussein, Serpil Tokdemir, Nitin Agarwal, and Samer Al-Khateeb (2018) 'Analyzing Disinformation and Crowd Manipulation Tactics on YouTube', 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM) . Available at https://ieeexplore.ieee.org/document/8508766 [accessed 14/04/2019]

[78] Brian Boland (2014) 'Organic Reach on Facebook: Your Questions Answered', Facebok Business. Available at https://www.facebook.com/business/news/Organic-Reach-on-Facebook [accessed 26/04/2019]; Roger McNamee (2019) Zucked: Waking Up to the Facebook Catastrophe , HarperCollins

[79] Roger McNamee (2019) Zucked: Waking Up to the Facebook Catastrophe , HarperCollins

[80] Tufekci (2016) p.68

[81] Edson C Tandoc, Jr and Julian Maitra (2017) 'News organizations' use of Native Videos on Facebook: Tweaking the journalistic field one algorithm change at a time', New Media & Society . Available at https://journals.sagepub.com/doi/abs/10.1177/1461444817702398 [accessed 14/04/2019]

[82] Suzanne Vranica and Jack Marshall (2016) 'Facebook Overestimated Key Video Metric for Two Years', The Wall Street Journal . Available at https://www.wsj.com/articles/facebook-overestimated-key-video-metric-for-two-years-1474586951 [accessed 14/04/2019]

[83] Adam Mosseri (2018) 'Bringing People Closer Together', Facebook Newsroom . Available at https://newsroom.fb.com/news/2018/01/news-feed-fyi-bringing-people-closer-together [accessed 14/04/2019]

[84] Laura Hazard Owen (2019) 'One year in, Facebook's big algorithm change has spurred an angry, Fox News-dominated — and very engaged! — News Feed', Nieman Lab . Available at https://www.niemanlab.org/2019/03/one-year-in-facebooks-big-algorithm-change-has-spurred-an-angry-fox-news-dominated-and-very-engaged-news-feed [accessed 14/04/2019]

[85] Nic Newman (2019) 'Journalism, Media, and Technology Trends and Predictions 2019', Reuters Institute Digital News Project , p.9. Available at https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2019-01/Newman_Predictions_2019_FINAL_2.pdf [accessed 14/04/2019]; Jim Waterson (2019) 'As Huffpost and BuzzFeed shed staff, has the digital content bubble burst?', The Guardian . Available at https://www.theguardian.com/media/2019/jan/24/as-huffpost-and-buzzfeed-shed-staff-has-the-digital-content-bubble-burst [accessed 14/04/2019]; Jill Abramson (2019) 'How Google and Facebook are Slowly Strangling Their Digital Offspring', Vanity Fair . Available at https://www.vanityfair.com/news/2019/01/how-google-and-facebook-are-strangling-their-digital-offspring [accessed 14/04/2019] ; see also Joshua A Braun and Jessica L Eklund (2019) 'Fake News, Real Money: Ad Tech Platforms, Profit-Driven Hoaxes, and the Business of Journalism', Digital Journalism , 7(1), pp.1-21. Available at https://www.tandfonline.com/doi/abs/10.1080/21670811.2018.1556314?journalCode=rdij20 [accessed 14/04/2019]

[86] Owen (2019)

[87] European Commission, Case AT.39740 - Google Search (Shopping) . Available at http://ec.europa.eu/competition/elojade/isef/case_details.cfm?proc_code=1_39740 [accessed 14/04/2019]

[88] Autorité de la concurrence, Decision 19-MC-01 of 3 January 2019 regarding a request for interim measures from Amadeus . Available at http://www.autoritedelaconcurrence.fr/user/standard.php?id_rub=697&id_article=3343&la [accessed 14/04/2019]

[89] Google (2010) 'Being bad to your customers is bad for business, Google Official Blog . https://googleblog.blogspot.com/2010/12/being-bad-to-your-customers-is-bad-for.html [accessed 14/04/2019]

[90] Owen (2019)

[91] Bundeskartellamt (2019) 'Bundeskartellamt prohibits Facebook from combining user data from different sources', Case B6-22/16 . Available at https://www.bundeskartellamt.de/SharedDocs/Meldung/EN/Pressemitteilungen/2019/07_02_2019_Facebook.html [accessed 14/04/2019]

[92] European Commission, Case M.8228 - Facebook / WhatsApp . Available at http://ec.europa.eu/competition/elojade/isef/case_details.cfm?proc_code=2_M_8228 [accessed 14/04/2019]

[93] Emma Goodman, Sharif Labo, Damian Tambini, and Martin Moore (2017) 'The new political campaigning', LSE Media Policy Project Series, Media Policy Brief 19 . Available at http://eprints.lse.ac.uk/71945 [accessed 14/04/2019]

[94] Robert M Bond, Christopher J Fariss, Jason J Jones, Adam D I Kramer, Cameron Marlow, Jamie E Settle, and James H Fowler (2012) 'A 60-million-person experiment in social influence and political mobilization', Nature , 489. Available at https://www.nature.com/articles/nature11421 [accessed 14/04/2019]; Jason J Jones, Robert M Bond, Eytan Bakshy, Dean Eckles, and James H Fowler (2017) 'Social influence and political mobilization: further evidence from a randomized experiment in the 2012 U.S. presidential election', PLoS One , 12(4). Available at https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0173851 [accessed 14/04/2019]

[95] Robert Epstein and Ronald E Robertso (2015) 'The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections', PNAS , 112(33). Available at https://www.pnas.org/content/112/33/E4512 [accessed 14/04/2019]

[96] Tambini (2018) p.286

[97] Daniel Kreiss & Shannon C McGregor (2018) 'Technology Firms Shape Political Communication: The Work of Microsoft, Facebook, Twitter, and Google With Campaigns During the 2016 U.S. Presidential Cycle', Political Communication , 35(2), pp.155-177. Available at https://www.tandfonline.com/doi/abs/10.1080/10584609.2017.1364814?journalCode=upcp20 [accessed 14/04/2019]

[98] Jonathan Zittrain (2014) 'Engineering an Election', Harvard Law Review Forum , 127. Available at https://harvardlawreview.org/2014/06/engineering-an-election [accessed 14/04/2019]

[99] Google (2012) SOPA/PIPA . Available at https://www.google.com/doodles/sopa-pipa [accessed 14/04/2019]; Zittrain (2014)

[100] Tambini, (2018) p.287

[101] Tambini (2018) p.287

[102] See, e.g., YouTube (2019) 'Continuing our work to improve recommendations on YouTube', YouTube Official Blog . Available at https://youtube.googleblog.com/2019/01/continuing-our-work-to-improve.html [accessed 14/04/2019] ; Monica Bickert (2019) 'Combatting Vaccine Misinformation', Facebook Newsroom . Available at https://newsroom.fb.com/news/2019/03/combatting-vaccine-misinformation [accessed 14/04/2019]; Mark Zuckeberg (2018) 'A Blueprint for Content Governance and Enforcement', Facebook . Available

at https://www.facebook.com/notes/mark-zuckerberg/a-blueprint-for-content-governance-and-enforcement/10156443129621634 [accessed 14/04/2019]; Mark Zuckerberg (2018) 'Preparing for Elections', Facebook . Available at https://www.facebook.com/notes/mark-zuckerberg/preparing-for-elections/10156300047606634 [accessed 14/04/2019]

[103] Luca Belli and Jamila Venturini (2016) 'Private Ordering and the rise of terms of service as cyber-regulation', Internet Policy Review , 5(4). Available at https://policyreview.info/articles/analysis/private-ordering-and-rise-terms-service-cyber-regulation [accessed 14/04/2019]; Lawrence Lessig (1999) Code and Other Laws of Cyberspace , Basic Books

[104] See Nick Clegg (2019) 'Charting a Course for an Oversight Board for Content Decisions', Facebook Newsroom . Available at https://newsroom.fb.com/news/2019/01/oversight-board [accessed 14/04/2019] ;Zuckerberg, 2018; Ezra Klein (2018) 'Mark Zuckerberg on Facebook's hardest year, and what comes next', Vox . Available at https://www.vox.com/2018/4/2/17185052/mark-zuckerberg-facebook-interview-fake-news-bots-cambridge [accessed 14/04/2019]

[105] Transposed into UK domestic law by The Electronic Commerce (EC Directive) Regulations 2002, SI 2002/2013 ('E-Commerce Regulations'). While data protection law is also highly relevant in this area, this paper focuses on platform liability protections.

[106] Technical Standards and Regulations Directive, art.1; see also E-Commerce Directive, Recital 18

[107] E-Commerce Directive, Section 4; Google France SARL and Google Inc. v Louis Vuitton (C-236/08), Google France SARL v Viaticum SA and Luteciel SARL (C-237/08) and Google France SARL v Centre national de recherche en relations humaines (CNRRH) SARL and Others (C-238/08) ECLI:EU:C:2010:159 (' Google France and Google ') at [112]

[108] E-Commerce Directive, Art.12

[109] E-Commerce Directive, Art.13

[110] E-Commerce Directive, Art.14

[111] E-Commerce Directive, Art.12, Art.13

[112] E-Commerce Directive, Art.14

[113] L'Oreal v eBay (C-324/09) at [111], [113]; Google France and Google (C-236/08) at [114]

[114] E-Commerce Directive, Art.14

[115] L'Oreal v eBay (C-324/09) at [113]; Google France and Google (C-236/08) at [114]

[116] E-Commerce Directive, Recital 42; Google France and Google (C-236/08) at [112-114]; L'Oréal SA and Others v eBay International AG and Others (C-324/09) ECLI:EU:C:2011:474 (' L'Oreal v eBay ') at [111]-[112]

[117] E-Commerce Directive, Recital 42

[118] E-Commerce Directive, Recital 42; Google France and Google (C-236/08) at [114]; L'Oreal v eBay (C-324/09) at [113]; note that as of December 2019 there is a case under consideration by the CJEU that may touch on this (LF v YouTube (C-682/18))

[119] L'Oreal v eBay (C-324/09) at [116]

[120] Google France and Google (C-236/08) at [114]

[121] E-Commerce Directive, Arts.12-14

[122] Google France and Google (C-236/08) at [116]

[123] See, by analogy, L'Oreal v eBay (C-324/09) at [116]

[124] Decision 7708/19 Reti Televisive Italiane SpA v Yahoo! Inc ; see also Alessandro La Rosa (2019) 'Ruling of the Italian Supreme Court in the Mediaset vs Yahoo! Case', Lexology . Available at https://www.lexology.com/library/detail.aspx?g=3390d24b-d614-409d-aed3-96f8c2d51e83 [accessed 25/04/2019]

[125] Decision 7708/19 Reti Televisive Italiane SpA v Yahoo! Inc at [4.3] - "le attivita di filtro, selezione, indicizzazione, organizzazione, catalogazione, aggregazione, valutazione, uso, modifica, estrazione o promozione dei contenuti, operate mediante una gestione imprenditoriale del servizio, come pure l'adozione di una tecnica di valutazione comportamentale degli utenti per aumentarne la fidelizzazione: condotte che abbiano, in sostanza, l'effetto di completare ed arricchire in modo no passivo la fruizione dei contenuti da parte di utenti indeterminati"

[126] England and Wales Cricket Board Limited and Another v Tixdaq Limited and Another [2016] EWHC 575 (Ch) at [167]-[170]

[127] Although, in the UK, service providers have qualified protections against defamation claims for content uploaded by users (Defamation Act 2013, s.5)

[128] Regulation (EU) 2019/1150 of the European Parliament and of the Council of 20 June 2019 on promoting fairness and transparency for business users of online intermediation services (Official Journal L 186, 11/7/2019, P. 57-79)

[129] Regulation (EU) 2019/1150, Recitals 7-8

[130] European Commission (2019) 'Platform-to-business trading practices'. Available at https://ec.europa.eu/digital-single-market/en/business-business-trading-practices [accessed 12/07/2019]

[131] Natali Helberger, Paddy Leerssen, and Max Van Drunen (2019) 'Germany proposes Europe's first diversity rules for social media platforms', LSE Media Policy Project . https://blogs.lse.ac.uk/mediapolicyproject/2019/05/29/germany-proposes-europes-first-diversity-rules-for-social-media-platforms/ [accessed 06/06/2019]

[132] Gordon Pennycook, Tyrone Cannon, and David G Rand (2018) 'Prior Exposure Increases Perceived Accuracy of Fake News', Journal of Experimental Psychology: General , 147(12). Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2958246 [accessed 14/04/2019]

[133] Christopher A Bail, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, M B Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky (2018) 'Exposure to opposing views on social media can increase political polarization', PNAS , 115(37). Available at https://www.pnas.org/content/115/37/9216 [accessed 14/04/2019]

[134] Michela Del Vicario, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi (2017) 'Modeling confirmation bias and polarisation', Scientific Reports , 7. Available at https://arxiv.org/abs/1607.00022 [accessed 14/04/2019]

[135] For example, Woods, Perrin, and Walsh (2019)

[136] For example, House of Lords (2018) 'Social Media and Online Platforms as Publishers: Debate on 11 January 2018', Library Briefing . Available at https://researchbriefings.parliament.uk/ResearchBriefing/Summary/LLN-2018-0003 [accessed 14/04/2019]

[137] For example, Department for Digital, Culture, Media & Sport (2019)

[138] See, for example, European Data Protection Board (2019) Guidelines 2/2019 on the processing of personal data under Article 6(1)(b) GDPR in the context of the provision of online services to data subjects . Available at https://edpb.europa.eu/our-work-tools/public-consultations/2019/guidelines-22019-processing-personal-data-under-article-61b_en# [accessed 14/04/2019]

[139] Indeed, as a form of personal data processing, data protection Supervisory Authorities are empowered to ban service providers (as data controllers) from recommending (as processing) in response to breaches of GDPR (see art.58)

[140] Facebook seems to have failed to meet non-discrimination obligations in relation to targeted advertising (Muhammad Ali, Piotr Sapiezynski, Miranda Boken, Aleksanra Korolova, Alan Mislove, and Aaron Rieke (2019) 'Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes', arXiv Preprints , arXiv:1904.02095v2. Available at https://arxiv.org/abs/1904.02095 [accessed 14/04/2019])

[141] E-Commerce Directive, Art.15

[142] L'Oreal v eBay (C-324/09) at [139]; Scarlet Extended SA v Société belge des auteurs, compositeurs et éditeurs SCRL (SABAM) (C-70/10) ECLI:EU:C:2011:771 at [36]-[40]; Belgische Vereniging van Auteurs, Componisten en Uitgevers CVBA (SABAM) v Netlog NV (C-360/10) ECLI:EU:C:2012:85 at [32]-[38]

[143] Aaron Smith (2018) 'Many Facebook users do not understand how the site's news feed works', Pew Research Center . Available at http://www.pewresearch.org/fact-tank/2018/09/05/many-facebook-users-dont-understand-how-the-sites-news-feed-works [accessed 14/04/2019]

[144] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, Christian Sandvig (2015) 'I always assumed that I wasn't really that close to [her]: Reasoning about invisible algorithms in the news feed', CHI '15 Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems . Available at https://dl.acm.org/citation.cfm?id=2702556 [accessed 14/04/2019]

[145] Emilee Rader and Rebecca Gray (2015) 'Understanding User Beliefs About Algorithmic Curation in the Facebook News Feed', CHI '15 Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems . Available at https://dl.acm.org/citation.cfm?id=2702174 [accessed 14/04/2019]; Bucher (2017)

[146] Lilian Edwards and Michael Veale (2017) 'Slave to the Algorithm? Why a 'Right to an Explanation' Is Probably Not the Remedy You Are Looking For', Duke Law & Technology Review , 16. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2972855## [accessed 14/04/2019]

[147] Adrian Weller (2017) 'Challenges for transparency', Paper presented at the 2017 ICML Workshop on Human Interpretability in Machine Learning (WHI 2017), Sydney . Available at https://arxiv.org/abs/1708.01870 [accessed 14/04/2019]

[148] Athanasios Andreou, Giridhari Venkatadri, Oana Goga, Krishna P Gummadi, Patrick Loiseau, and Alan Mislove (2018) 'Investigating ad transparency mechanisms in social media: A case study of Facebook's explanations', Proceedings of the Network and Distributed System Security Symposium (NDSS) . Available at http://www.eurecom.fr/en/publication/5414/detail/investigating-ad-transparency-mechanisms-in-social-media-a-case-study-of-facebook-s-explanations [accessed 14/04/2019]

[149] Delacroix (2019) p.29

[150] Jatinder Singh, Christopher Millard, Chris Reed, Jennifer Cobbe, and Jon Crowcroft (2018) 'Accountability in the IoT: Systems, Law & Ways Forward', Computer , 50(7), pp.54-65. Available at https://ssrn.com/abstract=3269792 [accessed 28/07/2019]

[151] https://algotransparency.org/ [accessed 14/04/2019]

[152] Ramya Sethuraman (2019) 'Why Am I Seeing This? We Have an Answer for You', Facebook Newsroom . Available at https://newsroom.fb.com/news/2019/03/why-am-i-seeing-this [accessed 14/04/2019]

[153] Susanne Barth and Menno D T de Jong (2017) 'The privacy paradox - Investigating discrepancies between expressed privacy concerns and actual online behavior - A systematic literature review', Telematics and Informatics , 34(7). Available at https://www.sciencedirect.com/science/article/pii/S0736585317302022 [accessed 14/04/2019]

[154] Christian Pieter Hoffmann, Christoph Lutz, and Giulia Ranzini (2016) 'Privacy cynicism: A new approach to the privacy paradox', Cyberpsychology: Journal of Psychosocial Research on Cyberspace , 10(4). Available at https://cyberpsychology.eu/article/view/6280/5888 [accessed 14/04/2019]

[155] René Wies (1995) 'Using a Classification of Management Policies for Policy Specification and Policy Transformation', in A S Sethi, Y Raynaud, and F Faure-Vincent (eds.), Integrated Network Management IV , IFIP — The International Federation for Information Processing, Springer. Available at https://link.springer.com/chapter/10.1007/978-0-387-34890-2_4 [accessed 14/04/2019]

[156] Vindu Goel (2018) 'India Curbs Power of Amazon and Walmart to Sell Products Online', The New York Times . Available at https://www.nytimes.com/2018/12/26/technology/india-amazon-walmart-online-retail.html [accessed 14/04/2019]

[157] Google Spain SL and Google Inc. v Agencia Española de Protección de Datos (AEPD) and Mario Costeja González (C-131/12) ECLI:EU:C:2014:317

[158] GDPR, Art.17

[159] GDPR, Art.5(1)(b)

[160] Christopher Marsden (2004) 'Co- and Self-regulation in European Media and Internet Sectors: The Results of Oxford University's Study www.selfregulation.info' in Christian Moller and Arnaud Amouroux (eds.) The media freedom internet cookbook , OSCE Representative on Freedom of the Media, p.80. Available at https://www.osce.org/fom/13844?download=true [accessed 14/04/2019]; see also Christopher Marsden (2011) Internet co-regulation: European law, regulatory governance and legitimacy in cyberspace , Cambridge University Press; and Michèle Finck (2017) 'Digital Co-Regulation: Designing a Supranational Legal Framework for the Platform Economy', LSE Legal Studies Working Paper No. 15/2017 . Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2990043 [accessed 14/04/2019]

[161] For example: Olivia Solon (2017) 'Google's bad week: YouTube loses millions as advertising row reaches US', The Guardian . Available at https://www.theguardian.com/technology/2017/mar/25/google-youtube-advertising-extremist-content-att-verizon [accessed 14/04/2019]; Patience Haggin and Suzanne Vranica (2019) 'Nestlé, McDonald's, Others Pull Ads From YouTube', The Wall Street Journal . Available at https://www.wsj.com/articles/nestle-mcdonalds-others-pull-ads-from-youtube-11550705643 [accessed 14/04/2019]

[162] For example, environmental regulation (James G Speight (2017) 'Environmental Regulations', in James G Speight, Environmental Organic Chemistry for Engineers , Butterworth-Heinemann)

[163] Diane Coyle (2018) 'Practical competition policy implications of digital platforms', Antitrust Law Journal , Forthcoming. Available at https://www.bennettinstitute.cam.ac.uk/media/uploads/files/Practical_competition_policy_tools_for_digital_platforms.pdf [accessed 14/04/2019]; Torsten Körber (2018) Is Knowledge (Market) Power? - On the Relationship Between Data Protection, 'Data Power' and Competition Law . Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3112232 [accessed 14/04/2019]

[164] For example, Lilian Edwards (2004) 'The Problem with Privacy', International Review of Law Computers & Technology , 18(3). Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1857536 [accessed 14/04/2019]

[165] For example, Lilian Edwards and Michael Veale (2018) 'Enslaving the Algorithm: From a 'Right to an Explanation' to a 'Right to Better Decisions'?', IEEE Security & Privacy , 16(3). Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3052831 [accessed 14/04/2019]; Sandra Wachter and Brent Mittelstadt (2018) 'A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI', Columbia Business Law Review , Forthcoming. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3248829 [accessed 14/04/2019]

[166] For example, Jatinder Singh, Christopher Millard, Chris Reed, Jennifer Cobbe and Jon Crowcroft(2018) 'Accountability in the IoT: Systems, Law & Ways Forward', IEEE Computer , 51(7). Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3269792 [accessed 14/04/2019]; Jat Singh, Jennifer Cobbe, and Chris Norval (2019) 'Decision Provenance: Harnessing Data Flow for Accountable Systems' , IEEE Access . Available at https://ieeexplore.ieee.org/document/8579125 [accessed 14/04/2019]

[167] House of Commons Digital, Culture, Media and Sport Committee (2019) pp.82-88