

III. Toward Constructive Optimisation: a new perspective on the regulation of recommender systems and the rights of users and society



Laurens Naudts, Natali Helberger, Marijn Sax, Michael Veale¹

- Introduction**..... 27
 - A. Recommenders and Society..... 28
 - B. Challenges in Realising Constructive Optimisation: Motivating our Approach..... 29
- PART 1: Normative Foundations**..... 31
 - 1.1. Relational Dynamics..... 31
 - 1.2 Self-Development and Self-Determination as Digital Values..... 32
 - 1.2.1 Self-determination..... 33
 - 1.2.2 Self-development..... 34
 - 1.3 Self-Determination and Self-Development as Social, Egalitarian and Structural Values..... 36
 - 1.4 The Supportive Function of the EU Charter of Fundamental Rights..... 37
- PART 2: The Stack Approach**..... 38
 - 2.1 The need for Constructive optimisation..... 38
 - 2.1 Contextualising the Stack..... 40
- PART 3: EU Regulation and The Stack**..... 41
 - 3.1 Toward Efficient Regulation: Toward efficient regulation along the optimisation stack..... 41
 - 3.2 Overarching Guidelines for Stack Governance..... 42
 - 3.3 Regulatory Recommendations Across the Stack..... 43

¹ For its conceptualization of self-development and self-determination as digital values, this study builds upon Naudts, Fair or Unfair Differentiation? Reconsidering the Concept of Equality for the Regulation of Algorithmically Guided Decision-Making (Doctoral Dissertation, KU Leuven: Leuven, 2023). In reference to: Iris Marion Young, Justice and the Politics of Difference (Princeton University Press 1990) and Iris Marion Young, Inclusion and Democracy (Oxford University Press 2002)

- 3.3.1 Business to Consumer Interface (Hardware).....44
- 3.3.2 Business to Consumer Interface (Software).....45
- 3.3.3 Functionality50
- 3.4 Engine52
 - 3.4.1 Alternative engines.....52
- 3.5 Input Data55
 - 3.5.1 Input Data (Content)56
 - 3.5.2 Input data (users and environments).....58
- 3.6 Business-to-Business Interface.....61
- 3.7 Connectivity Infrastructure.....62
- 3.8 Whole-Stack Governance.....62
 - 3.8.1 Operations and Management.....62
 - 3.8.2 Accountability Groups.....64
 - 3.8.3 Individuals and Communities.....65
- Conclusion67**

Introduction

How should we regulate systems designed to optimise digital environments and interactions?

One needs to develop at least two critical perspectives to answer such a question. First, relative to what normative standards should optimisation be held? Second, how should regulation understand the tools of optimisation, such as ‘recommender systems’? This study develops an approach to both questions and integrates the corresponding perspectives into one answer.

The study is divided into three main parts. In **Part 1** a normative framework – centred around the values of self-development and self-determination – is elaborated as an interpretational resource to understand better how optimisation can be meaningful. When it comes to recommender systems, there is a need to move beyond naïve approaches, which implicitly assume that ‘the recommender system’ is an identifiable, discrete ‘unit’ that can be addressed and regulated as such. Instead, we propose to conceptualise and evaluate recommender systems through a so-called “stack approach”. This is the purpose of **Part 2**. The envisaged “stack approach” embraces the insight that beyond the surface interface level, recommender services are the result of different interactions, operations and layers, that are both social and technical in nature — software, hardware, infrastructure, organisational, design principles, and so on. All these parts work in concert to, ultimately, create particular tools, interfaces, and functionalities. Finally, **Part 3** combines the normative framework of Part 1 and the stack approach of Part 2 for a critical analysis of the current approach to the regulation of recommender systems under the DSA, and for developing constructive suggestions of how to better account for the legitimate interests of users and society. Recommender systems should be regulated addressing *every layer of the stack*. Put simply, analysing and regulating the recommender system is not (only) about analysing and regulating the actual recommender engine, i.e., the software systems designed to fulfil optimisation logics, or the interface people interact with. The net should be cast wider. Optimisation goals determined by management, KPIs determined by business departments, performance reviews, hiring practices, data collection and analysis practices, iterative software design philosophies, UX/UI design choices, data training models, and so on, should all be incorporated into the bread and butter of recommender system regulation.

This study, then, combines a more realistic, helpful approach to recommender systems as socio-technical artefacts with an original theoretical perspective on the normative standards we should hold optimisation systems to. In this report, we formulate a set of overarching recommendations that could guide future regulatory amendments. In an upcoming update and annex to this report, we will take up this exercise ourselves, and demonstrate how our model can be translated into concrete regulatory provisions. At the same time, we offer the stack approach as a toolkit to the reader: a starting point for reflection toward a more healthy and fair digital eco-system. In this context, it should be noted from the outset that the more realistic stack approach can be as enlightening as it can be overwhelming. The benefit of the approach is that it allows for a very wide, structural approach that cuts across the *entire* recommender value-chain or stack to show how a wide range of EU legislation can be used to regulate various elements of this ecosystem. The resulting analysis can, at the same time, also lead to what feels like a rather fragmented story – at least in terms of presentation. To further add to this enlightening complexity, the stack approach allows one to address separate layers of the stack individually, but one can also show how several layers (can) interact with one

another in the regulatory context, or how ‘whole stack provisions’ address the entire stack. In short, the stack’s *analytical* modularity allows for a very all-encompassing mosaic approach that can address several analytical levels at the same time. Its inherent complexity is a feature, not a bug. This should be kept in mind when reading this exploratory study.

A. Recommenders and Society

Physical and digital settings are increasingly subject to systems that attempt to optimise humans and their interactions.² In various public and private domains, systems affect the content people see, whether advertisements or commercial product offerings, audio-visual entertainment, news media, potential professional and personal connections, etc.³ In short, recommender systems have become integral to structuring the digital society. In this function, they actively co-mediate people’s social and economic affordances.⁴ Their underpinning logic however, is one of ‘capture’ and ‘traps’, where through their interaction with these systems, individuals’ behaviours are codified and computed, and their actions and attention are steered into certain business logics that can be difficult to escape from.⁵ Given their ubiquity, recommender systems can, in theory, help with information overload — helping maximise user freedom, filtering content that is more catered to the needs and desires of the recipient, thereby reducing the time they would otherwise lose when confronted with cognitive overload.⁶ This was the dream of early proponents of these systems, where ‘adaptive hypermedia’ would allow users to achieve their goals more easily.⁷ On a societal level, recommender systems used by news websites can incorporate diversity metrics to promote voices that are otherwise left unseen and unheard. Yet, those same data-driven techniques can also be used in the opposite direction. Individuals and groups can find their social and economic practices captured and subject to manipulation.⁸ Already marginalised groups can be rendered even more invisible.⁹

Recommender system’s alignment with democratic norms and values, fundamental rights, freedoms and interests, greatly depends upon the optimisation strategies followed within them. Yet, as Kulynych and others warn, when they are “developed to capture and manipulate

-
- 2 Bogdan Kulynych and others, ‘POTs: Protective Optimization Technologies’, *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (ACM 2020) <<http://dl.acm.org/doi/10.1145/3351095.3372853>> accessed 29 January 2020.
 - 3 Max van Druenen, Brahim Zarouali and Natali Helberger, ‘Recommenders You Can Rely on: A Legal and Empirical Perspective on the Transparency and Control Individuals Require to Trust News Personalisation’ (2022) 13 JIPITEC <<https://www.jipitec.eu/issues/jipitec-13-3-2022/5562>>. Paddy Leerssen, ‘Seeing What Others Are Seeing: Studies in the Regulation of Transparency for Social Media Recommender Systems.’ (2023).
 - 4 Iason Gabriel, ‘Towards a Theory of Justice for Artificial Intelligence’ (2022) 151 *Daedalus* 12.
 - 5 Nick Seaver, ‘Captivating Algorithms: Recommender Systems as Traps’ (2019) 24 *Journal of Material Culture* 421; Philip E Agre, ‘Surveillance and Capture: Two Models of Privacy’ (1994) 10 *The Information Society* 101.
 - 6 See John Danaher, ‘Freedom in an Age of Algocracy’ in Shannon Vallor (ed), John Danaher, *The Oxford Handbook of Philosophy of Technology* (Oxford University Press 2022) <<https://oxfordhandbooks.com/view/10.1093/oxfordhb/9780190851187.001.0001/oxfordhb-9780190851187-e-16>> accessed 4 May 2022.
 - 7 Peter Brusilovsky, ‘Adaptive Hypermedia’ (2001) 11 *User Modeling and User-Adapted Interaction* 87.
 - 8 See for instance: Marijn Sax, *Between Empowerment and Manipulation: The Ethics and Regulation of for-Profit Health Apps* (2021) <<https://dare.uva.nl/search?identifier=52225d37-e7e1-4883-9dab-a3fd3a063d8>> accessed 18 September 2023.
 - 9 See for example: Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York University Press 2018) <<http://www.degruyter.com/document/doi/10.18574/9781479833641/html>> accessed 8 December 2021; Catherine D’Ignazio and Lauren F Klein, *Data Feminism* (The MIT Press 2020) <<https://mitpress.mit.edu/books/data-feminism>>; Sanne Vrijenhoek and others, ‘Recommenders with a Mission: Assessing Diversity in News Recommendations’, *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval* (ACM 2021) <<https://dl.acm.org/doi/10.1145/3406522.3446019>> accessed 7 September 2023.

behaviour and environments for the extraction of value, [they introduce] broader risks and harms for users and environments beyond the outcome of a single algorithm within that system.”¹⁰ While these systems are commonly called ‘personalised’, this term is misleading. All too often, ‘personalisation’ masks self-serving optimisation. They rarely treat individuals as individuals, focussing on their contribution towards aggregate outcomes. They rarely allow users to set their own ‘personal’ goals. Likewise, the systemic threats recommenders pose to democratic and societal interests, interactions and structures are too often overlooked. Instead, the goals are those set by firms with decisional power over the implementation and control over the infrastructural, data and knowledge resources needed for their design. Sometimes, these firms are traditional economic actors, but often, they are intermediaries or ‘platforms’, where both sides of a transaction are relatively powerless compared to those setting and shaping the rules of engagement.¹¹

The importance of socio-technical systems optimising people and environments means that they should be designed to align with the citizen-consumer and the democratic and social values the EU promotes, from the outset and from the cradle to the grave.¹² Such ambitions are also reflected in the European Declaration on Digital Rights and Principles, which promotes human-centricity and freedom of choice, solidarity, inclusion and participation in the democratic process as key commitments in the EU’s digital transformation.¹³

B. Challenges in Realising Constructive Optimisation: Motivating our Approach.

To realise the lofty ambitions of a just digital future, several challenges must be overcome. First, among the various values we could choose from as leading technological advancements, which and whose to prioritise? This is why in Part 1 of this study; we offer a theoretical framework which serves as an indispensable background to navigating questions on values. The digital ecosystem is comprised of multiple actors and groups who each might pursue differing and competing interests and values. Our current information society is further characterised by significant asymmetries in power over (physical and digital) infrastructures, design choices, expertise, knowledge of consumers, and data creation processes. One of the main current tools of optimisation online, the recommender system, therefore, often operates behind a

¹⁰ Bogdan Kulynych e.a., ‘POTs: Protective Optimization Technologies’, in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* ’20: Conference on Fairness, Accountability, and Transparency, Barcelona Spain: ACM, 2020)*, 177–88, <https://doi.org/10.1145/3351095.3372853.moral>, social, and political impact that digital systems have on populations through solutions that can be applied by service providers. Fairness frameworks do so, in part, by mapping these problems to a narrow definition and assuming the service providers can be trusted to deploy countermeasures. Not surprisingly, these decisions limit fairness frameworks’ ability to capture a variety of harms caused by systems. We characterize fairness limitations using concepts from requirements engineering and from social sciences. We show that the focus on algorithms’ inputs and outputs misses harms that arise from systems interacting with the world; that the focus on bias and discrimination omits broader harms on populations and their environments; and that relying on service providers excludes scenarios where they are not cooperative or intentionally adversarial. We propose Protective Optimization Technologies (POTs)

¹¹ Julie E Cohen, *Between Truth and Power: The Legal Constructions of Informational Capitalism* (Oxford University Press 2019).

¹² Natali Helberger, ‘On the Democratic Role of News Recommenders’ (2019) 7 *Digital Journalism* 993; Marijn Sax, ‘Algorithmic News Diversity and Democratic Theory: Adding Agonism to the Mix’ (2022) 10 *Digital Journalism* 1650.

¹³ European Declaration on Digital Rights and Principles for the Digital Decade, see also: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/europes-digital-decade-digital-targets-2030_en#digital-rights-and-principles; https://ec.europa.eu/commission/presscorner/detail/en/ip_22_7683

veil of opacity.¹⁴ Consequently, it has become increasingly difficult to assess whether these socio-technical systems pursue merely self-serving interests or comprehensively consider people's fundamental rights, democratic values and social well-being. It has been difficult even to understand the contribution of these systems to outcomes both online and offline.

Moreover, in choosing specific values over others, we must also consider the inherent trade-offs that often need to be made to safeguard the interests of the individual on the one hand and those of (social) collectives on the other. Maximising individual user preferences might come at the cost of content diversity.¹⁵ In certain domains, like news provision, optimising for corporate goals might deny certain social groups the visibility needed to organise themselves politically. Likewise, pursuing short-term goals, like personal relevance-based metrics, might ultimately have a detrimental long-term impact. For example, it might diminish people's opportunity to encounter unexpected voices or perspectives and develop new interests. Likewise, if consumption is encouraged, what is the impact thereof on the environment? Tensions between interests will undoubtedly arise, and answers on how they can be resolved depend upon one's normative and political outlook. Still, when choosing the values we want to promote within the digital environment, we can take stock of these complexities. More specifically, we need values that enable us to capture, reflect on, and address these conflicts as they emerge; values that recognise how people's wellbeing is not determined in light of one particular preference, need or desire, but shaped by a multitude of interweaving factors. Only then will we be able to steer optimisation systems and recommender systems to work in favour of — or at least not to the significant detriment of — the broader health and welfare of the information ecosystem.

Second, to ensure a healthy digital environment, we must address its complexity. This is what Part 2 of the study is focused on. Optimisation systems have 'many hands' involved and are not as easy to steer and regulate as some other phenomena.¹⁶ We need to understand how recommender systems interact with their surroundings to assess their possibilities and limitations, risks and benefits, the conditions under which they can do good, and the conditions that impede that goal. Optimisation systems are already soaked in EU regulation, but we have to ask whether these initiatives hit the right targets, work together in concert, and achieve their aims. Existing technology laws, such as the Digital Services Act or the GDPR, (should) exert influence on design, development and deployment, but in practice, may not be aligned or refined enough to do so. Opportunities that exist under existing law may not be taken full advantage of. In mapping how recommenders, and those who control their value-chain, interact with and affect their surroundings, we can assess whether current laws adequately capture the (systemic) risks these systems threaten to impose onto citizens, and in case of regulatory failure, propose legislative recommendations to mitigate future harm.

¹⁴ Naudts (n 1).

¹⁵ Helberger (n 12).

¹⁶ A Feder Cooper and others, 'Accountability in an Algorithmic Society: Relationality, Responsibility, and Robustness in Machine Learning', *2022 ACM Conference on Fairness, Accountability, and Transparency* (ACM 2022) <<https://dl.acm.org/doi/10.1145/3531146.3533150>> accessed 23 June 2022; Helen Nissenbaum, 'Accountability in a Computerized Society' (1996) 2 *Science and Engineering Ethics* 25; Jennifer Cobbe, Michael Veale and Jatinder Singh, 'Understanding Accountability in Algorithmic Supply Chains', *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Association for Computing Machinery 2023) <<https://dl.acm.org/doi/10.1145/3593013.3594073>> accessed 14 June 2023. Natali Helberger, Jo Pierson and Thomas Poell, 'Governing Online Platforms: From Contested to Cooperative Responsibility' (2018) 34 *The Information Society* 1.

An important connection between Part 1 and Part 2 is the thematisation and capturing of *relational dynamics*. As one can see directly below, in Part 1 we formulate a theoretical framework which explicitly acknowledges the *relationality* inherent in our current technological predicament. Various stakeholders interact with each other, citizens interact with technology and do so differently in different contexts. Likewise, the values and principles we might want to see realised, can play out differently in each of these constellations. Part 2 also works with this relationality by introducing the stack approach which is aimed precisely at analytically disentangling the relationships that exist within and between the many layers of the stack. By laying bare these relational dynamics, the role the (constitution of) technology plays in societal relational dynamics can be better analysed. It's relationality all the way down.

This study addresses these challenges via the following strategy. In Part 1, we first argue that recommender systems can only be made sense of when viewed in light of society's structural, institutional and relational dynamics. Second, we posit self-development and self-determination as leading values that should inform recommender systems' development, integration and regulation. In Part 2 we lay out a framework to better understand the structural and social dynamics that give meaning to and are influenced by recommender systems. To this end, we propose visualising these socio-technical systems through a *stack* of interrelated decision-making moments, infrastructural capacities, steps and processes. In part III, we analyse the current regulatory approach to recommender systems through the lens of constructive, accountable and societally aligned optimisation across the entire stack. Throughout this analysis, we make recommendations for tweaks, initiatives and improvements.

PART 1: Normative Foundations

1.1. Relational Dynamics

For citizens and consumers, recommender systems perform an active mediating role in navigating the digital society. In attributing recommender systems as active co-mediators, we acknowledge a certain fluidity to their functioning. Their interventions are dynamic, and this dynamicity occurs alongside three broad dimensions.

First, recommender systems are informed by an interplay of interpersonal, institutional, and socio-technical relationships. For example, recommender systems operate on data collected from a wider population. Those data might reflect the prejudiced beliefs and stereotypes held by society's members. The norms and values held by individuals, communities, professions and industries, might equally inform the choices during the design, development and deployment of recommender systems. Certain industries and fields might be characterised by a vehement belief in techno solutionism as the answer to all of societies woes.¹⁷ At the same time, recommender systems must comply with the EU legislator's principles, norms and rules. And within these broader structures, actors in recommender systems take a wide variety of roles and perspectives.¹⁸

¹⁷ Abeba Birhane, 'Algorithmic Injustice: A Relational Ethics Approach', *Patterns (New York, N.Y.)* 2, nr. 2 (12 February 2021): 100205, <https://doi.org/10.1016/j.patter.2021.100205>.{\i}Patterns (New York, N.Y.

¹⁸ For instance, as *data controllers* they are obliged to incorporate data protection by design standards. Likewise, recommender systems are rightfully expected not to act in discriminatory ways. As *providers of AI systems* (and depending on the nature of the final AI Act) they may be obliged to treat their systems as (high-)risk. As *contractual*

Second, and upon deployment, citizens and consumers interact with recommender systems. As recommendations are provided, they simultaneously adapt to their user's preferences. This also means that users, through their interaction and the data they provide, contribute to the performance and improvement of a recommender system.¹⁹ Yet, dependent on the goals imbued within the system, the consumer is expected, or steered, to behave and act in specific ways. Individuals on the other side of optimisation systems give a lot, including the labour of navigating them and the externalities they create, but have little say in how they function, or mechanisms to have them aligned with their own interests.

Third, recommenders maintain an interactive relationship with the individual citizen consumer and draw from and co-shape the (social) position of consumers vis-à-vis other consumers. Citizens are grouped together and thus positioned in a relationship with each other based on their habits, preferences or other monitorable behaviours. Depending on the context and circumstances, the latter could impact, both positively and negatively, the social positioning of individuals and the groups they are a member of. For instance, the offering to young men of clothing embroidered with slogans affirming stereotyped gender patterns could damage women's social position. They mix the economic and the expressive and produce and reproduce socioeconomic structures. Regulating optimisation is a hazardous business because the potential goals are varied, contextual, dynamic and political. Yet to leave them insufficiently steered and regulated leaves these significant choices in solely private hands.

1.2 Self-Development and Self-Determination as Digital Values²⁰

Drawing from Iris Marion Young's political philosophy, we put self-determination and self-development as the values a just society should promote to safeguard people's dignity and equal moral worth. Due to optimisation systems contribution to the structuring of society, these principles and values should be respected throughout the digital value chain if they are to be realised. From a regulatory perspective, the objective should be twofold. Positively, regulators should maximise people's capacity for self-development and self-determination. Negatively, regulators should remove social and institutional barriers that can negatively impact people's access to self-development and self-determination. Those regulatory efforts can take various forms and target different actors. They can comprise technical design strategies, organisational requirements, procedural contestation measures, impact assessments, data audits, etc. Importantly, however, there is no one-size-fits-all solution.

Before analysing the regulatory efforts to govern recommender systems across the stack, we first establish the interconnected notions of self-development and self-determination and

parties in the AI supply chain, they may have interconnected obligations with cloud and infrastructural providers. Yet as *intermediaries* under the Digital Services Act, they may have both platform-related obligations while benefiting from certain immunities for the actions they take or content they host.

19 Balázs Bodó and others, 'Interested in Diversity' (2019) 7 *Digital Journalism* 206; Masoud Mansoury and others, 'Feedback Loop and Bias Amplification in Recommender Systems' (arXiv, 25 July 2020) <<http://arxiv.org/abs/2007.13019>> accessed 28 July 2023.

20 Drawing from Young's politics of difference and democratic theory, the ideation and application of both notions within the digital environment used in this study were first introduced in Naudts, *Fair or Unfair Differentiation? Reconsidering the Concept of Equality for the Regulation of Algorithmically Guided Decision-Making* (Doctoral Dissertation, KU Leuven: Leuven, 2023), Chapter. 6, Socio-Relational Conceptualisations of Equality and Algorithmically Guided Decision-Making. Iris Marion Young, *Justice and the Politics of Difference* (Princeton University Press 1990) and Iris Marion Young, *Inclusion and Democracy* (Oxford University Press 2002)

the social ambitions these values seek to realise. Second, we want to position these values as social rather than purely individual.

1.2.1 Self-determination

The value of self-determination concerns people's ability "to participate in determining one's action and the condition of one's action."²¹ To exercise this type of control, people need access to social and material resources which enable them to exercise choice.²² Conversely, people do not have self-determination when others can arbitrarily interfere with their exercise of choice. Where individuals and groups are subject to external optimisation strategies, these practices and the systems involved, are perhaps best understood as affecting the mental, physical, social and material resources people need access to in exercising choice and control over the conditions that govern their lives.²³

If designed carefully, there are situations where recommender systems could positively complement people's ability to exercise choice. For one, Danaher notes, they can help citizens 'identify and select among options that might (or might not) be conducive to [their] goals' because they 'filter choices and reduce the feeling of being overwhelmed.'²⁴ Yet without appropriate governance mechanisms, recommender systems can turn the online environment into a space of arbitrary or unchecked control.²⁵

Following a republican perspective, the current position held by operators might enable 'algorithmic domination'.²⁶ Drawing from citizens' data, behaviour or other monitorable actions, recommender systems limit, replace or create the options, content and goods people see and access in favour of the goals of the optimisation.²⁷ This control over choice is problematic because it is *outsourced* without the citizen-consumers' deliberative engagement. Indeed, citizens often have little to no say, influence, or insight into what most recommender systems have been optimised for. Instead, the operators of recommender systems have been placed in a position where they can control the choices offered to citizens under their sole discretion and according to their preferences.²⁸

Regulatory efforts should attempt to make the digital environment more favourable, inclusive and participatory as to include the interests of citizens, social groups, and society at large. Those efforts can take two strategies. First, regulators should enhance and broaden citizens' resources to exercise deliberate control over recommender systems.²⁹ And, where consumers agree to have their choices curated by others, they should understand the rules of engagement.

²¹ Iris Marion Young, *Inclusion and Democracy* (Oxford University Press 2002) 32 <<http://www.oxfordscholarship.com/view/10.1093/0198297556.001.0001/acprof-9780198297550>> accessed 7 October 2020.

²² Philip Pettit, *On The People's Terms: A Republican Theory and Model of Democracy* (Cambridge University Press 2012) 45.

²³ Naudts (n 1).

²⁴ Danaher (n 6).

²⁵ E Gräf, 'When Automated Profiling Threatens Our Freedom': (2017) 3 *European Data Protection Law Review* 441, 450.

²⁶ Philip Pettit, *Republicanism: A Theory of Freedom and Government* (Oxford University Press 1999) <<https://www.oxfordscholarship.com/view/10.1093/0198296428.001.0001/acprof-9780198296423>> accessed 29 February 2020; Pettit (n 22); Gräf (n 25); *ibid*; Danaher (n 6).

²⁷ Gräf (n 25) 450.

²⁸ The notion of operators should be broadly understood, referring to a wide variety of actors within the technology stack, including those who provide the infrastructural resources needed to maintain and sustain their functioning.

²⁹ See also: Gräf (n 25).

People must understand the conditions under which curation takes place, have the opportunity to contest certain decisions that impact their choices, and, where possible, they themselves or their representatives should be able to participate in the design of recommender systems. It does bear explicit mentioning that the ideation of citizens as deliberate actors should not be interpreted as favouring disclosure and consent as the chosen regulatory solution. Quite the opposite. Among others, and as we show below, recommender systems are opaque, not simply because they are technically arcane but because their outcomes are the result of distributed decisions across an entire technological stack. There is no easy transparency fix. Even if citizens have been given some degree of control over how they navigate their digitally mediated environment, their doing so is often conditional upon someone else's goodwill: infrastructural and knowledge asymmetries characterise people's digital living space, those with the power to determine how and when recommender systems will be designed and deployed often can unilaterally decide how and when people's choices will be interfered with. Given these external constraints, it has become increasingly difficult for citizens to understand exactly what they would consent to. That does not mean that disclosure and consent cannot be appropriate under certain circumstances. In (complex) digital environments however, they can be too easily abused by operators to reallocate the responsibility in the incurrence of injustice and harm back to the citizen.

A better interpretation of the notion of self-determination is to further democratise the regulation of recommender systems, rather than have the rules of, and goals pursued by, optimisation unilaterally decided upon and imposed onto them. In other words, citizens should be given the possibility to co-determine how their living sphere will be mediated through digital technologies, rather than opt-in, through consent, to an environment, the conditions of its structuring, has already been predetermined for them. The effective ability for citizens to determine their actions and the conditions of their actions, must therefore be realised through a wide range of measures, which moreover, must be interpreted in line with the below-mentioned value of self-development.

Second, regulators should minimise and constrain operators in their ability to (deliberately) undermine the interests of citizens in pursuit of self-serving ones. These two regulatory strategies could moreover be combined by promoting inclusivity and active participation as part of regulatory governance strategies, whether they pertain to the broader infrastructures and architectures in which systems are embedded, the technical design choices and parameters that guide a recommender system's optimisation goals or the law-making process, including its enforcement, itself. In this context, it is also valuable for citizens to have access to meaningful alternatives. Where the digital environment is dominated by a select few actors, people's ability to choose a different service provider that offers other conditions for content recommendations is limited, and their freedom to exercise choice becomes further reduced. Moreover, in doing so, we can maximise the positive externalities recommender systems provide, such as their ability to filter choice to limit informational glut.³⁰

1.2.2 Self-development

The value of self-development can be interpreted as people's ability to 'learn and use satisfying and expansive skills in socially recognised settings, and enable them to play and communicate

30 Danaher (n 6).

with others or express their feelings and perspectives on social life in contexts where others can listen.³¹ This value has been linked to the capability approach.³² Under this approach, capabilities are understood as the substantial freedoms or opportunities people have to achieve alternative beings and doings.³³ Though people vary in their conception of the good life, in a pluralist society, people should be free to pursue what they deem valuable. And people only have a real opportunity to pursue their goals when they have access to the appropriate emotional, psychological, social, institutional and material supportive mechanisms to do so. Though material resources are an important benefactor, people's capacity to convert their opportunities also depends on their cultural and social positioning and recognition as equals.³⁴

Once again, socio-technical systems act as an external (environmental) condition that can either limit or enhance people's ability to flourish. These systems help structure who is offered what type of content and when, and therefore also what content will remain invisible to whom. For example, whereas a price differentiation system can determine who has access to certain goods at more favourable prices (material resources), news recommenders can determine whose voices will be heard, and whose will be silenced (cultural and social). In turn, the choices underlying recommender systems affect people's interpretation and navigation of the digital society.

For example, imagine a mass-used video-on-demand platform whose recommendations promote programming that reinforces gendered stereotypes of women. In perpetuating prejudiced and generalised assumptions regarding the social position of women, such a recommender system could actively undermine women's capacity to 'develop their personal abilities, pursue their professional careers and/or make choices about their lives.'³⁵ Likewise, if news recommender systems never offer news content related to the struggles faced by marginalised or vulnerable communities, the average citizen might remain insensitive to their needs. In the worst case, those communities are rendered invisible, and their members unable to express themselves in socially recognised settings, such as public political discourse.

Self-development is at stake both from the substance of optimisation but also is affected by the conditions under which this optimisation can be controlled. Where controls around these systems exist but are in practice unusable or opaque, those resources to contest undesirable decisions might as well not be there. Such conditions require a consideration of the entire optimisation process, from top to bottom, to ensure individual citizen consumers, as well as the social groups they are a member of, are recognised and heard.

31 Young (n 1) 31–32; Iris Marion Young, *Justice and the Politics of Difference* (Princeton University Press 1990).

32 Amartya Sen, 'Equality of What?' (1979); Amartya Sen, *The Idea of Justice* (Belknap Press: Harvard University Press 2009); Martha C Nussbaum, *Creating Capabilities: The Human Development Approach* (The Belknap Press of Harvard University Press 2011).

33 Sen (n 32); Nussbaum (n 32); Ingrid Robeyns, 'The Capability Approach' in Edward N Zalta (ed), *The Stanford Encyclopedia of Philosophy* (Winter 2016, Metaphysics Research Lab, Stanford University 2016) <<https://plato.stanford.edu/archives/win2016/entries/capability-approach/>> accessed 18 December 2019.

34 Young (n 31); Nancy Fraser, 'From Redistribution to Recognition? Dilemmas of Justice in a "Post-Socialist" Age' [1995] *New Left Review* 68; Amy Allen, 'Power and the Politics of Difference: Oppression, Empowerment, and Transnational Justice' (2008) 23 *Hypatia* 156.

35 <https://www.ohchr.org/en/women/gender-stereotyping>

1.3 Self-Determination and Self-Development as Social, Egalitarian and Structural Values

Self-development and self-determination should not be viewed solely as individual values but as values that represent, and are informed by, collective, social, egalitarian, and structural dynamics.

As the above examples illustrate, recommender systems can mirror and reinforce the (historical) prejudice, stereotypes and stigma faced by marginalised or otherwise vulnerable communities and their members. These types of injustice act as social, economic and institutional barriers that limit people's ability to be seen, heard, and recognised (self-development), which, in turn, undermines their capacity to exercise deliberative choice over the conditions that govern their life (self-determination). However, tackling prejudiced recommenders cannot be resolved through individual or isolated interventions alone. For instance, when confronted with a biased recommender system, an individual right to opt-out from personalisation does little to address the problem at its core. Without structural interventions, the normalisation of prejudiced worldviews through technology persists. Because individuals can face disadvantages due to their membership in social groups, mechanisms should be available to enhance the collective ability of these groups to contest and evaluate recommender systems.

At the same time, recommender systems perform group-level operations under the guise of personalisation.³⁶ The information relied on is aggregated on a population level, and the targets optimised toward are similarly aggregated. Optimisation systems convey information about individuals as group members rather than individuals as individuals (e.g. *people* who watch *x* might like *y*).³⁷ Consequently, in structuring the world as we see it, recommender systems do not affect single individuals but groups of individuals. Moreover, as these groups are typically not stable enough to talk about as delineated collectives or for them to socialise and mobilise, it renders the negotiation, representation and dialogue of tech governance more difficult.

Of course, people receive content from various sources, and recommender systems continuously perform actions that might limit or impede people's self-development and self-determination. Hence, their influence might not seem as impactful when viewed in isolation. The point we make, however, is to acknowledge that these systems, due to their integration and widespread use, exert control over people's conditions to participate in social and economic life and that the origins of this influence might manifest itself on a collective level first, before ultimately harming individuals as members of these collectives. Hence, the realisation of individual self-development and self-determination, as well as the protection of group interests, needs to be performed through individual-, collective-, and societal-level interventions.³⁸

In identifying these collective dynamics, we might also see tensions arise between individual and group-level interests. For instance, in specific contexts, such as news curation, it could

³⁶ Salomé Viljoen, 'A Relational Theory of Data Governance' [2021] the Yale Law Journal 82.

³⁷ See also: Anton Vedder, 'KDD: The Challenge to Individualism' (1999) 1 Ethics and Information Technology 275; Anton Vedder and Laurens Naudts, 'Accountability for the Use of Algorithms in a Big Data Environment' (2017) 31 International Review of Law, Computers & Technology 206; Viljoen (n 36).

³⁸ Given the widespread interest in and popularity of foundation models, the problem of harm arising on a collective level will likely persist. Indeed, foundation models "are not built for a specific *context* or conditions of use, and their openness and ease of control allow for unprecedented *scale* of use." See also: Natali Helberger and Nicholas Diakopoulos, 'ChatGPT and the AI Act' (2023) 12 Internet Policy Review <<https://policyreview.info/essay/chatgpt-and-ai-act>> accessed 25 May 2023.

be beneficial to increase the visibility of marginalised groups, even if this would contradict the personal interests of individual users.

In trying to resolve such tensions, it is best to understand self-development and self-determination as egalitarian values. As observed by Young, these values “assume the equal moral worth of all persons, and thus justice requires their promotion for *everyone*.”³⁹ Yet, in a society characterised by social and economic inequality, we should acknowledge that not every person currently has *equal* access to the means for self-development and self-determination. To assure self-development and self-determination for everyone, however, the regulation of recommender systems should be informed by the broader social and institutional infrastructures in which these systems are embedded and with which they interact. In this context, the structuring function of recommender systems is often informed by background conditions of injustice, as evidenced by examples in which they reinforce historical prejudice.⁴⁰ Regulators and operators should, therefore, show awareness of the disadvantage people might experience because of group membership, paying particular attention to the social and economic position held by marginalised or otherwise vulnerable communities. To positively contribute to a healthy digital living environment, recommender systems should be harnessed to dismantle, rather than perpetuate, structural inequalities.⁴¹ This also means sufficient – and often more – corporate, procedural, cultural and institutional resources should be made available to ensure their interests are represented and prior injustice rectified. Though the social recognition and economic integration of marginalised communities can only be realised through concerted and society-wide efforts among them, technology regulation should be no exception. These efforts are needed to ensure every person is able to deliberate and participate on *equal footing* as individuals or community members, either alone or collectively through their representation via social interest groups, in creating (ex-ante) and evaluating (ex-post) recommender systems.

1.4 The Supportive Function of the EU Charter of Fundamental Rights

The effective enjoyment of fundamental rights is a necessary external institutional condition that enables EU citizens to exercise self-development and self-determination. At the same time, fundamental rights concern areas of life in which people should be seen and heard and be able to exercise control as to how they can enjoy these rights. In other words, whereas people require fundamental rights to effectuate self-development and self-determination, self-development and self-determination are needed to fully enjoy their fundamental freedoms. In sum, fundamental rights are an inviolable North Star in the recommender systems’ design, deployment and evaluation. Still, as Kas notes in Chapter V of this anthology, the Charter’s position as an institutional bulwark in the protection of citizen consumers can be further strengthened. Indeed, more efforts will be needed to realise the Charter’s potential in realising citizen consumers’ interests to be free from digitally generated harm and injustice through fundamental rights.⁴²

39 *Justice and the Politics of Difference* (Princeton University Press 1990) 37.

40 Annette Zimmermann and Chad Lee-Stronach, ‘Proceed with Caution’ [2021] *Canadian Journal of Philosophy* 1.

41 Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police and Punish the Poor* (Macmillan Publishers 2018).

42 Betül Kas, Ensuring Digital Fairness in EU Consumer Law through Fundamental Rights: Is the EU Charter Fit for Purpose?

PART 2: The Stack Approach

2.1 The need for Constructive optimisation

The logics of optimisation are so deeply entwined with the fabric of society, there will seemingly always be some role for them in digital society going forward. They are part of the infrastructure of the digital public sphere, and as such have an important societal dimension. Yet society has few mechanisms to hold them to account, and to align them so that their goals reflect a plurality of interests, rather than a corporate actor with outlandish decisional power. Given their influence is likely to endure, regulating them is more realistic than removing them.

The societal dimension is in parts already recognised in the law, and the DSA more specifically, though primarily through the lens of the potential (systemic) risks to consumers, society and societal values and freedoms⁴³ and the need to protect consumers from arbitrary and opaque decisions that influence the ability to determine their actions.⁴⁴ Having said so, a mere risk-based approach fails to account for the critical role of recommenders as a means to make a positive and meaningful contribution to the realisation of users' rights to self-determination and self-development. Instead, it is imperative to create the external conditions through which people can flourish not only map the risks but use technology to dismantle barriers.

We, therefore, argue that users should not only have a legitimate interest to be protected from risks that recommenders pose. There is also a need to acknowledge and respect the freedom to use recommender systems as a tool to understand and navigate the digital information economy, to advance their rights to self-development and self-determination using recommender systems as means of discovery and learning, to increase their visibility within society, to pursue the information goals they deem valuable, etc. A core deficit with optimization today is that it is difficult for any actor to hold specific optimization logics to account, to politicise them, to reject, refute or refuse them, or to steer them. Specific optimization logics are baked into the services consumers use and are to be accepted by consumers as a fact of life. Concretely, this means that the regulation of recommender systems cannot be left to the sole discretion of operators of recommender systems to decide the goals to optimise the recommendation algorithm.

Unfortunately, within the current climate, users cannot just stand up and leave to a different provider that has their political values. Users, then, have a legitimate interest in that a recommender system is optimised in a way that is inherently useful and meaningful to them. To do so requires a space of contestation to steer these values, dynamically and continuously. Constructive and accountable optimisation is a starting point for this. If optimisation systems can be made more societally responsive, not (just) at the level of the response of individuals within them but at the level of their aims, purposes and governance, then we may protect against harms while enabling positive navigation of digital society.

Constructive optimisation therefore not only refers to operators' ability to justify and defend the normative choices they have made. Constructive optimisation mandates the ex-ante

⁴³ Art. 35 and 36 DSA.

⁴⁴ Art. 27 DSA. See also: Jennifer Cobbe and Jatinder Singh, 'Regulating Recommending: Motivations, Considerations, and Principles' (2019) 10 European Journal of Law and Technology <<https://ejlt.org/index.php/ejlt/article/view/686>> accessed 28 July 2023.

possibility for end-users, civil society groups, regulators and others to participate in the processes through which those choices are made. Likewise, it should be possible to scrutinise and contest those choices ex-post.

To render optimisation systems more socially responsive, we need to reconfigure our understanding and representation of these technologies, considering the above-mentioned relational dynamics.

Operators and regulators would surely like nothing more than to tug out the ‘algorithm’, examine and assess it, and hold those responsible for its function to account. Recommendation and optimisation systems are not single pieces of software. There may well be underlying technologies but there are also layers of (latent) social norms and values, business rules, meetings, varied logics, oversight, manual intervention, alpha and beta tests, new interfacing products, features and initiatives, content moderation and compliance. There is not “one algorithm” to locate. Instead, recommenders must be understood as socio-technical, systems that have been designed, but whose functioning is socially embedded and constituted.⁴⁵

To capture this complexity, we propose the metaphor of the optimisation stack as a more comprehensive imagination of the dynamic socio-economic technology structures that recommender systems are.

Thinking in terms of a stack does not get us all the way to grappling with the ‘structure of economic relationships that data systems support’, but it does unveil some of the components of these systems, and some of the functions of these data systems. Other areas of law and regulation — and politics in general — is necessary to consider reshaping economic systems.

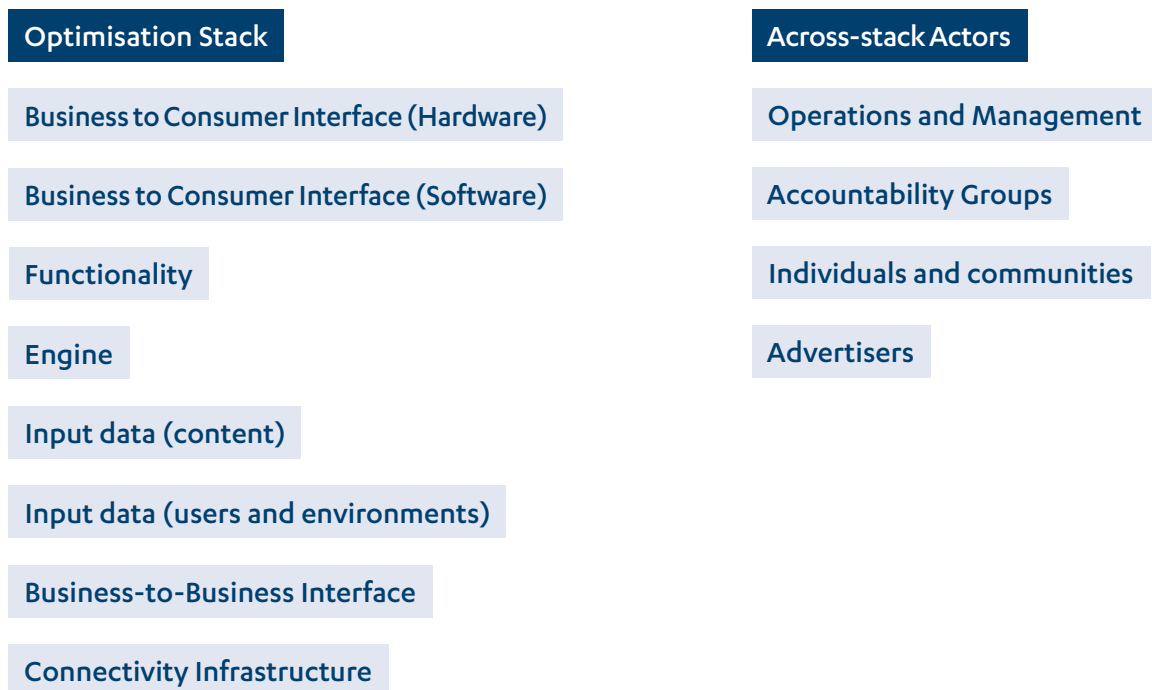
The stack we propose has the following components.

- **Business to Consumer Interface (Hardware)** — *the material way individuals interact with a service, which constraints the possibilities for certain governance interventions*
- **Business to Consumer Interface (Software)** — *the software interface that is delivered through hardware interface(s), potentially more dynamic, adjustable, and individualised*
- **Functionality** — *more abstract capabilities of computing systems than interfaces, in this layer we find tasks that computing systems are designed to achieve for users, providers and others*
- **Engine** — *software systems designed to fulfil optimisation logics, drawing on data to provide functionality, interfaces and more*
- **Input data (content)** — *expressive forms of information that we associate with speech norms*
- **Input data (users and environments)** — *descriptive or predictive forms of information relating to users*
- **Business-to-Business Interface** — *the ways that businesses interact with a service and in turn with consumers, which facilitate certain interactions and business practices over others*
- **Connectivity Infrastructure** — *generic Internet technologies.*

⁴⁵ Nick Seaver, *Computing Taste: Algorithms and the Makers of Music Recommendation* (University of Chicago Press 2022).

Sitting parallel to the stack we have several actors and agencies which inform many parts of the optimisation stack simultaneously — they are not part of it, but governing, steering and supporting them is key to understanding governance within the stack.

- **Operations and Management** — *Organisational functions within businesses intended to oversee all or part of the stack. Without capacity and links across the stack in this function, governance mechanisms would not be implementable*
- **Accountability Groups** — *Inclusive of research groups, journalists, civil society, standardisation bodies and regulators. Without capacity in this function, there would be a very limited audience for constructive accountability, transparency and contestation*
- **Individuals and (social) communities** – *Whether they find representation through the abovementioned accountability groups or not*
- **Advertisers.**



2.1 Contextualising the Stack

The stack metaphor recalls the layered architecture of the Internet. The Internet as we experience it is architecturally constructed from different abstract levels of technology, which both interact with each other to produce the end results we see, but are also conceptually isolated, insofar as the lowest levels of the Internet, which transfer packets end-to-end across the network, do so without regard to the applications being run on them, like the Web or e-mail. Scholars have emphasised that there are many perspectives to look at the Internet from, including this architectural approach, and a synthesis of perspectives is required for effective governance.⁴⁶

⁴⁶ William Lehr and others, 'Whither the Public Internet?' (2019) 9 Journal of Information Policy 1.

We argue that the same treatment needs to be given to recommender systems in order to ensure technology functions robustly. This follows from other work considering digital power and European governance of technologies through a stack model undertaken by the Ada Lovelace Institute.⁴⁷ One reason that a stack is a useful way to conceive an optimisation system is that individual and societal outcomes of interest from recommender systems stem from the interaction of many parts of the system. Not all parts of the system are controlled by the same actor, and over time, the actors who govern the stack and the functions within it might change.

We can also use the metaphor (and practice) of optimisation to describe what is going on in the digital consumer sector.⁴⁸ In other troubling sectors which have been accused of significant illegal behaviour, such as online display advertising, ‘decentralisation’ has made regulation more difficult. Different actors run different parts of these services. Where they are all controlled by a single actor, there may be hope to regulate through a particular ‘choke-point’ — such as an all-encompassing platform.⁴⁹ But e-commerce in Europe is subject to many players, who are increasingly part of complex algorithmic supply chains.⁵⁰ These set-ups are economically designed to separate liability from the actor extracting value from a platform, as the history of intermediary liability law tells us.⁵¹ A website or a seller themselves is going to be enmeshed in a complex platform ecosystem, where their developers may feel powerless when faced with the changing services of a large platform they are integrated with.⁵² Users too might even be attributed some responsibility for the governance of recommender and platform systems.⁵³

PART 3: EU Regulation and The Stack

3.1 Toward Efficient Regulation: Toward efficient regulation along the optimisation stack

In Part I, we put forward self-determination and self-development as critical, social, egalitarian, and structural values that can be interpreted as informed by, and in response to, the relational dynamics that constitute the information society. In Part II, the optimisation stack was developed to account for the dynamicity of the digital environment. The conceptual foundation in the former, and technological ideation in the latter of this study, inform one another. Both

47 Valentina Pavel and others, ‘Rethinking Data and Rebalancing Digital Power’ (Ada Lovelace Institute, 2022) <<https://www.adalovelaceinstitute.org/wp-content/uploads/2022/11/Ada-Lovelace-Institute-Rethinking-data-and-rebalancing-digital-power-FINAL.pdf>>.

48 Bogdan Kulynych and others, ‘POTs: Protective Optimization Technologies’, *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (ACM 2020).

49 Jack L Goldsmith and Tim Wu, *Who Controls the Internet? Illusions of a Borderless World* (Oxford University Press 2006).

50 Jennifer Cobbe, Michael Veale and Jatinder Singh, ‘Understanding Accountability in Algorithmic Supply Chains’, *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Association for Computing Machinery 2023).

51 Lilian Edwards, ‘“With Great Power Comes Great Responsibility?”: The Rise of Platform Liability’ in Lilian Edwards (ed), *Law, Policy, and the Internet* (Hart Publishing 2019).

52 Tania Bucher, ‘Objects of Intense Feeling: The Case of the Twitter API’ (2013) 3 *Computational Culture: A Journal of Software Studies*.

53 Natali Helberger, Jo Pierson and Thomas Poell, ‘Governing Online Platforms: From Contested to Cooperative Responsibility’ (2018) 34 *The Information Society* 1.

are integral to our endeavour to critically interrogate the regulation of recommender systems, which is the objective of Part 3.

Regulation cannot take aim at a single part of the system. For example, regulation that tries to moderate content misses the underlying recommender engine. Regulation aimed at the underlying logics of the recommender engine ignore the design dynamics which lead to certain patterns of input. Regulation which focuses on certain design dynamics of input might ignore the role or potential of communities to flag, steer, use or repurpose these systems. The stack offers a unique vantage point, and toolkit for others, which facilitates the identification of relevant legislation in a piecemeal manner, whereby each layer within the stack might be subject to different rules or even regulatory frameworks. Once we surmise which laws can apply to which layer, we can identify which provisions could positively contribute to self-development and self-determination. As the stack visualises interaction between layers, once laws have been mapped, the visualisation can help analyse how laws interact, including where they complement or conflict. Moreover, upon disentangling the stack, it becomes evident how the regulation of the online environment becomes a long-term project:

There is no one-size-fits-all regulatory solution for turning the digital environment into a living space conducive to human flourishing. Instead, technology governance is a step-by-step process dependent upon various institutions' and individuals' efforts and contributions, aided by multiple interacting laws.

Turning to the values of self-development and self-determination, we have shown how their realisation cannot be tied solely to metrics directed toward personal or individual relevance (as a typical commercial optimization goal). The right to constructive optimisation, as grounded in the right to self-determination and self-development, includes the need to take into account the legitimate interests of users, both as individuals and as members of social groups and non-social collectives. Self-determination and self-development can moreover not be realised by focusing on short-term goals and choices. Their enjoyment for all is dependent upon a well-functioning and healthy democratic society imbued by a strong respect for fundamental rights. Hence, due consideration must be given to the ways in which novel technologies can interfere with the social and institutional structures in which citizens navigate, as well as the physical, mental, social and economic capacities citizens require, to determine their actions, and the conditions of those actions on the one hand, and expand, develop, express and communicate their experiences and perspectives on the other hand.

3.2 Overarching Guidelines for Stack Governance

Drawing from our analysis thus far, and before we start our evaluation of applicable EU legal frameworks, we propose a set of overarching standards aimed to facilitate users to take active part in the governance of socio-technical eco-systems in ways that are seen and heard. More specifically, regulation should be positively aimed in assuring that those subjugated to recommender systems can:

- understand the rules of engagement, including how systems function, for which purposes they have been optimised, and the consequences such optimisation strategies entail on the content users see
- have an actual and actionable say in the optimisation goals pursued within digital ecosystems

- exercise meaningful choice and voice, which requires the presence of alternative options, both in relation to a particular recommender system's functioning and in relation to other operators, including service-providers and platforms
- be included, represented and having one's voice heard during the ideation, design, deployment and evaluation of recommender systems, meaning they should have access to participation and contestation mechanisms from cradle to the grave.

Moreover, to ensure equal enjoyment of self-development and self-determination, and equal participation in attempts to democratise recommender systems, particular attention must be paid to marginalised or otherwise vulnerabilised communities. If not, the digital ecosystem will become an additional barrier to break down in the fight against the structural inequalities and injustice they already face. In this context, constructive optimisation also pertains to the interests of those who may not directly interact with a recommender or are (structurally) underrepresented and excluded because recommendation algorithms are not only a tool for users to discover information but also a means for non-users and other parties affected by the algorithmically mediated choices to be discovered.⁵⁴ Through their recommendation logics, recommender systems determine if and who gets seen and heard under which conditions, but also: who remains invisible and unheard. Likewise, the optimisation stack should be governed by strong labour protection and respect the rights of data workers involved in data production processes.⁵⁵

After having thus defined the cornerstones of constructive optimisation as a balance between individual and societal interests whereby each person has access to the conditions needed to exercise self-determination and self-developments, we will in the next step undertake a critical analysis of the current legal approach to regulating recommender systems in the DSA and adjunct laws.

3.3 Regulatory Recommendations Across the Stack

In the following section, we outline the stack piece-by-piece. In each section, we characterise this aspect of the optimisation stack, including why it may be challenging to govern. In addition, we indicate any regulatory provisions from the current or proposed European digital *acquis* which relates to this section (or the absence thereof). We then make proposals within this section for rights that relate to this aspect of the stack.

This exercise is explicitly exploratory in nature. We aim to show how a stack approach affords one with a new, different perspective to approach questions on recommender system regulation. The overall aim is thus not to be exhaustive, but rather to show the reader how a stack approach to recommender systems allows one to draw on the entire (current or proposed) European digital *acquis* to address many different aspects related to optimisation strategies such as recommender systems but also where potential intervention points are to realise more

⁵⁴ Philip Napoli and Sheea Sybblis, 'Access to Audiences as a First Amendment Right: Its Relevance and Implications for Electronic Media Policy' [2008] McGannon Center Working Paper Series <https://research.library.fordham.edu/mcgannon_working_papers/6>.

⁵⁵ See also the work of: Milagros Miceli e.a., 'Documenting Data Production Processes: A Participatory Approach for Data Work' (arXiv, 9 augustus 2022), <http://arxiv.org/abs/2207.04958>; Milagros Miceli en Julian Posada, 'The Data-Production Dispositif' (arXiv, 24 mei 2022), <http://arxiv.org/abs/2205.11963>; Milagros Miceli, 'Whose Truth? Power, Labor, and the Production of Ground-Truth Data' (Technische Universität Berlin, 2022), <https://depositonce.tu-berlin.de/handle/11303/19464.9> augustus 2022

accountable optimisation. As we already highlighted in the introduction, this exploratory stack approach can feel very mosaic, but that is more of a feature than a bug – the approach is meant to open up space for applying the European digital acquis to recommender systems. So, if anything, Part 3 should inspire the reader to get creative and use the stack as a model to embrace a wider, more dynamic perspective on recommender system regulation.

3.3.1 Business to Consumer Interface (Hardware)

Summary: the material way individuals interact with a service, which constraints the possibilities for certain governance interventions.

Hardware is not typically steered directly by legislation, and usually should not be. Yet, choices in hardware design can affect citizens in how they exercise and enjoy their rights. Where hardware decisions risk limiting people’s ability to co-govern recommender systems, and these negative externalities cannot be effectively designed around, operators must correct this imbalance through other means.

Hardware is difficult to govern directly as it often supports many services and is released or designed prior to the development of services or platforms on it. To link them would be undesirable as it would restrict openness and competition. Legislators have proceeded warily with functional mandates relating to consumer hardware, with legislative action on harmonised chargers for mobile devices taking many years to come to fruition.⁵⁶

Hardware decisions could positively contribute to people’s navigation of the digital environment. For instance, visually-impaired consumers benefit from certain hardware modalities, such as audio or haptics, to experience their surroundings and to exercise rights. At the same time, hardware might also constrain the effectiveness of rights or obligations that can be placed on services that use it. For example, obligations on prior transparency before data collection,⁵⁷ the use of recommender systems,⁵⁸ or attached to pieces of content⁵⁹ can be more difficult to implement on interfaces like smart speakers that rely on auditory cues only. The compromises service-providers need to make as a consequence of hardware choices could thus leave a legislative gap. Ideally then, the obligations of actors in the optimisation stack should alter depending on the modalities through which they deliver their service.

⁵⁶ Directive (EU) 2022/2380 of the European Parliament and of the Council of 23 November 2022 amending Directive 2014/53/EU on the harmonisation of the laws of the Member States relating to the making available on the market of radio equipment OJ L 315/30.

⁵⁷ e.g. GDPR, art 13.

⁵⁸ e.g. DSA, art 29.

⁵⁹ e.g. AIA Parliament Draft, art 52 [deepfake transparency]; DSA, art 31 [on traders].

Recommendation: Citizen-consumers' ability to exercise agency over recommender systems should not be negatively impacted due to their reliance on devices that have limited hardware support, but instead accommodated for. Among others, firms should be obliged to vary the form of transparency and ability to exercise agency according to the modalities their services are available on. Moreover, if there are true and hard limits to active and passive user rights that can be fulfilled, then other obligations or design strategies may need to be heightened and strengthened. An active reflection on such limitations should be promoted, for instance, through the performance of (periodic) risk assessments and ex-post monitoring obligations.

3.3.2 Business to Consumer Interface (Software)

Summary: The software interface that is delivered through hardware interface(s), which are potentially more dynamic, adjustable, and individualised in nature.

Interfaces are central to the optimisation stack, and are studied heavily in academic fields such as *interaction design*. Several existing provisions in the EU acquis relate directly and indirectly to business-to-consumer interfaces. In this section, we focus on the transparency of recommender system adjustments, the prohibition of manipulative design interfaces and the need to offer trader information.

3.3.2.1 Recommender System Adjustments and Transparency

At the core of the DSA's recommender provisions is the provision on recommender system transparency. This provision concentrates on the information presented to consumers. For one, providers of platforms that use recommender systems should "set out in their terms and conditions, in plain and intelligible language, the main parameters used in their recommender systems, as well as any options for the recipients of the service to modify or influence those main parameters."⁶⁰ This information is intended to give recipients insight into the why certain recommendations are made, and include the criteria most significant in determining the output suggested, and the reasons behind their relative importance.⁶¹ Article 27(3) DSA (Recommender system transparency) concludes with an obligation for platforms that have multiple options for the configuration of recommender systems to make available a 'functionality that allows the recipient of the service to select and to modify at any time their preferred option'. Such functionality must be 'directly and easily accessible from the specific section of the online platform's online interface where the information is being prioritised'. These guarantees however hardly reach further down the stack into the operational and technical level of the service itself.

There is no general obligation for online platforms to offer multiple or alternative options to end-users concerning the configuration of recommender systems they are subject to. Indeed, only where the platform itself has chosen to make alternative options available, should end-users have the ability to easily and directly navigate between those alternatives. In other words,

⁶⁰ Article 27, § 1 DSA.

⁶¹ Article 27, § 2 DSA.

meaningful, including alternative, choice is not viewed as an entitlement, but rather a privilege granted by platforms.

In this context, it is relevant to more profoundly consider what an ‘option’, and choice among options, would constitute. In practice, recommender systems are often deployed as ‘hybrid’ systems comprising many recommender systems. Mixed recommender systems stack the output of many recommenders — producing the content from a system perhaps trained on your behavioural data first, then following it by content that might be popular in your geographic region, and when data on that runs out, content that might be popular throughout the world.⁶² Ensemble systems rank and weigh many recommenders and combine the results. These options are not necessarily surfaced to users (typically), even though they typically do exist in the backend. Individuals could be provided with just the most popular content in their geography, for example, even though knowledge concerning their likely preferences behaviour-wise could be given.

In ensuring granularity of choice over options, lessons could be drawn from data protection law. It could be argued that in order to comply with data protection law, systems *have to be* created to be decomposable: they must be able to run with limited profiling or personal data, otherwise their design is contrary to Article 25 GDPR (Data protection by design and by default).⁶³

Recommendation: Regulators should take an expansive reading of Article 27 DSA and make it mandatory to offer consumers a meaningful choice through which they can realise their rights to self-determination and self-development, through among others, the choice among alternatives as a default.

In addition to the DSA, other interface-based information rights concerning the use of metrics or “main parameters” as part of recommender and ranking systems have been introduced. New additions in FB2BPR Article 5, UCPD Article 7(4a) and CRD Article 6a(a) requires users to be able to access the ‘main parameters’ of a ranking system when they search or conclude contracts with traders in various online settings. The purpose of these provisions when aimed at businesses is supposedly to ‘improve predictability’ and ‘improve the presentation’ of goods and services,⁶⁴ whereas for individuals the motivations are not explicit in the recitals.

While the term ‘main parameters’ includes ‘any general criteria, processes, specific signals incorporated into algorithms or other adjustment or demotion mechanisms used in connection with the ranking’, the instruments are clear that such parameters need not be customised per user.⁶⁵ A core problem however is that many recommender systems are tailored by user, including by having their weights and parameters significantly adjusted by a user’s activity within a certain session.⁶⁶ In this case, the value of the generic notion ‘main parameters’ seems unclear, and in any case, highly technologically specific. This also means that aggregate analysis or reasoning based on these main parameters may be flawed, as there is potential for unwanted dynamics

⁶² Kim Falk, *Practical Recommender Systems* (Manning 2019), ch 12.

⁶³ Case C-252/21 *Meta Platforms Ireland* ECLI:EU:C:2023:537.

⁶⁴ FB2BPR, recital 24.

⁶⁵ Omnibus Directive, recitals 22–23.

⁶⁶ See eg Mounita Bhattacharya and Sudarshan Lamkhede, ‘Augmenting Netflix Search with In-Session Adapted Recommendations’ (arXiv, 5 June 2022) <<http://arxiv.org/abs/2206.02254>> accessed 28 September 2023.

such as discrimination or manipulation on an individual level based on variations across recommender manifestations, or the exclusion of voices in society.

A second challenge with ‘main parameters’ is that where the input data (see below) to recommender systems is abstract in nature, such as originating from telemetry data including gyroscopes, compasses, wireless connections, clicks and touches of devices, or in virtual reality applications, even eye gaze, it is very difficult to explain systems *in terms* of these data.⁶⁷ Explanation facilities for machine learning and similar systems struggle to communicate with humans when the building blocks of the data underlying a system are abstract. It is easy to imagine a system explaining that ‘you received this advert because you are listed as enjoying holidays in Malta’. However, when a system is recommending this on the basis of telemetry data, there is no higher-level human concept that serves as an obvious mapping. As humans we appear to analyse the world in terms of these higher concepts — or at least we often *post hoc* reason that we do in our own heads. Machines don’t need to. As such, ‘main parameters can make little sense.

The above challenges broadly relate to the implementation of a right to ‘know your digital alter ego’, something that has been suggested in previous work.⁶⁸ Collectively, this would enable citizens to understand with whom they share their digital alter ego. Such a right requires us to reflect on what it would be to *know* such an alter ego given that it may be complex, extremely multi-faceted, dynamic and based on very abstract concepts with little human-interpretable meaning. European law in this area often attempts to square this circle by requiring complex information in plain and intelligible language, or similar. This is not an easy circle to square though, as a lot is lost by simplification — so much that it might undermine the faithfulness of an explanation as a whole. Presumably the purpose of transparency provisions is not to give comfort, but to allow individuals to benefit from and instrumentalise knowledge, which requires that knowledge can be mapped both onto their experiences and the functioning of the underlying systems. Transparency, action, and the building of understanding should all be interlinked.⁶⁹ This would seem to require a turn to more interactive and action-focused forms of transparency. ‘Main parameters’ simply does not do this work — the requirement is for systems to be created in more playful, exploratory ways, that facilitate users understanding, customising and altering them.

67 Lilian Edwards and Michael Veale, ‘Slave to the Algorithm? Why a “Right to an Explanation” Is Probably Not the Remedy You Are Looking For’ (2017) 16 Duke Law & Technology Review 18, 59.

68 https://www.beuc.eu/sites/default/files/publications/BEUC-X-2023-020_Consultation_paper_REFIT_consumer_law_digital_fairness.pdf

69 See eg Cynthia Rudin, ‘Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead’ (2019) 1 Nature Machine Intelligence 206. See eg Motahhare Eslami and others, ‘First I “Like” It, Then I Hide It: Folk Theories of Social Feeds’, Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (ACM 2016); Nadia Karizat and others, ‘Algorithmic Folk Theories and Identity: How TikTok Users Co-Produce Knowledge of Identity and Engage in Algorithmic Resistance’ (2021) 5 Proceedings of the ACM on Human-Computer Interaction 305:1. Sarah Inman and David Ribes, ‘“Beautiful Seams”: Strategic Revelations and Concealments’, Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Association for Computing Machinery 2019). Elias Storms, Oscar Alvarado and Luciana Monteiro-Krebs, ‘“Transparency Is Meant for Control” and Vice Versa: Learning from Co-Designing and Evaluating Algorithmic News Recommenders’ (2022) 6 Proceedings of the ACM on Human-Computer Interaction 405:1.

Recommendation: The informative utility of giving insight into a system’s ‘main parameters’ is flawed given the dynamic nature of recommendation and optimisation, and the abstract nature of underlying data. Information provided should enable users to understand the logic of the recommender system, the normative choices made therein, as well as the consequences of their own choices exercised in their interaction with recommender systems. This includes individual consequences (what users will see or not see) but also how companies took into account wider societal interests, such as the health of the public sphere, the interests of marginalised communities, the ecological footprint of technologies used and initiatives to protect workers’ rights. Moreover, to assess the trustworthiness of the systems they partake in, they should also be told by whom those choices were made and to which end. The end-users should thus understand how the aforementioned areas affect what the end-user is and is not recommended and with whom they share their so-called digital alter ego. Transparency obligations should (and can) be adapted to account for the dynamism of the optimisation stack. To be able to act upon information, information should be clearly presented, rather than hidden away in opaque and complex terms and conditions.

Recommendation: The Commission should explore participatory design obligations, potentially in collaboration with research capacities or projects, where platforms have to make transparency interactive, meaningful, and linked to action as to accommodate the needs and interests of end-users. Such efforts should moreover pay sufficient attention and include marginalised, or otherwise vulnerable, communities and interest groups.

3.3.2.2 Prohibitions on Manipulative Interfaces

Article 25 DSA (Online interface design and organisation) prohibits online platform providers from designing interfaces in a way that ‘manipulates the recipients of their service or in a way that otherwise materially distorts or impairs the ability of the recipients of their service to make free and informed decisions’, except where they relate to practices governed by the Unfair Commercial Practices Directive or the GDPR. The Commission can issue guidelines on how this article functions. This is the so-called ‘dark patterns’ provision. Paragraph 2 gives three examples, including presenting certain choices more prominently than others, using pop-ups to interfere with users’ choices and making the procedure for terminating a service more complex than subscribing to it. The limited scope of the provision aside (excluding the GDPR and the UCPD), Article 25 DSA merely tackles the *presentation* of choices and not the actual choices themselves. Those choices can still derive from data-driven and dynamically adjustable operations that identify and commercially exploit people’s preferences. Studies on dark patterns have indicated that it is both the presentation (e.g., how many layers down in a pop-up a decision-point is) and the content (e.g. is there a reject all button) that materially affects individuals’ choices.⁷⁰

Elsewhere in this report we discussed the lack of conceptual clarity in combination with the prerogative of platforms and private standardisation bodies to define the notion unilaterally as a fundamental problem of Article 25.⁷¹ The stack perspective adds another critique: Arti-

⁷⁰ Midas Nouwens and others, ‘Dark Patterns after the GDPR: Scraping Consent Pop-Ups and Demonstrating Their Influence’, *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2020)* (ACM 2020).

⁷¹ See Chapter II above - Helberger N, Sax M, Digital Vulnerability and Manipulation in the Emerging Digital Framework.

Article 25 DSA is no alternative for reaching further down the systems architecture and providing guidance for considering consumer interests (and vulnerabilities), such as regarding the way platform architectures are designed to collect data in the first place (**input data** level), how decisions about consumer data are made in the **operations and management capacity**, or which parameters the model is optimised for on the **engine** level.

Other EU law also contains constraints on manipulative interfaces which are slightly more general in nature and potentially reach beyond the interface level and deeper down into the technology stack. Article 5 UCPD prohibits a commercial practice, understood broadly, if it ‘materially distorts or is likely to materially distort the economic behaviour with regard to the product of the average consumer whom it reaches or to whom it is addressed, or of the average member of the group when a commercial practice is directed to a particular group of consumers.’

Article 5(1)(a) Draft AIA [EP] proposes that it is prohibited to place on the market, put into service or use ‘an AI system that deploys subliminal techniques beyond a person’s consciousness or purposefully manipulative or deceptive techniques, with the objective to or the effect of materially distorting a person’s or a group of persons’ behaviour by appreciably impairing the person’s ability to make an informed decision, thereby causing the person to take a decision that that person would not have otherwise taken in a manner that causes or is likely to cause that person, another person or group of persons significant harm’. This reaches further down the stack somewhat as the AI system itself is considered as the underlying engine, however the recitals also limit the full connection across the stack. Recital 16 notes that the ‘intention to distort the behaviour may not be presumed if the distortion results from factors external to the AI system which are outside of the control of the provider or the user, such as factors that may not be reasonably foreseen and mitigated by the provider or the deployer of the AI system.’ One could argue that this recital excludes “intent” where the distortion can be linked to the input data for example. Furthermore, excluding factors that are “external to the AI system”, appears to eliminate the need to duly consider the pervasive influence social conditions will have on the functioning of a system. To illustrate this tension, we can ask whether a system optimising for profit, where user or business input data makes it in practice act in a manipulative manner, would be considered as being “intended to manipulate.”

Consumer law, the AIA and data protection law together represent a regime where firms are obliged to develop interfaces in *modular* ways, with manipulative components identified and prohibited. As it stands, manipulative systems which optimise to extract value from consumers are the norm, and the default. Personalised recommendations should be a staged, directed, opt-in system, where individuals are involved in choosing the reasoning behind presenting information, and to make their goals more explicit. This does not mean that we return to a world of search; simply that when individuals are presented with adaptive interfaces, much like they ‘sort’ products in a search, they can choose options that reflect their aims and preferences. Perhaps they would like to see products other people with similar purchasing histories, or some other data they bring, purchased? Perhaps they would like to see products that are more like other products when they browse? Perhaps they would like interfaces based on popularity with their geography, or other demographics. Ultimately, firms’ ability to optimise for pure profit or clicks should be reduced — not removed, but placed as a subsidiary factor to an overarching logic that is chosen by a user.

Recommendation: Manipulative tactics undermine the opportunity for end-users to exert deliberate control over their digital environment (self-determination). Moreover, when end-users are manipulated to act in accordance with the operator's interests, they are also denied the opportunity to have their desires and interests heard and recognised (self-development). Whereas adaptive interferences often operate on a logic that is optimised for an aggregate concept such as profit or attention, and while there may be a role for including such goals in optimisation, they should be subsidiary to functional goals actively chosen by users. These considerations should inform the interpretation of Article 25 DSA, as well as the Commission when they issue guidelines on the application of Article 25 DSA on specific practices.

3.3.2.3 Product and Trader Information

Some further provisions interplay with the interface, effectively attempting to ensure that traders' information is effectively provided to consumers in the presence of intermediaries. Article 3(5) UCPD obliges online intermediation services to ensure that the identity of a business user providing goods is clearly available. Article 31 DSA (Compliance by design) obliges online platforms that allow consumers to contract with traders to ensure that interfaces are designed in such a manner which allows traders to comply with information provision obligations. Article 32(2) DSA (Right to information) requires providers to alter their online interfaces to inform consumers of illegal products if they cannot inform all concerned individuals directly. However, these are lopsided provisions. While platforms benefit from designing interactive, dynamic interfaces to deliver their business aims, regulatory goals cannot benefit from these tactics. If platforms can design compelling experiences for their own business ends, why can they not be obliged to design compelling experiences for regulatory ends? If platforms have information on what people might purchase, why should they not be obliged to 'advertise' to those same consumers, through optimization systems, that certain products they may have purchased on or off the service have been deemed to be illegal by a regulatory authority? This is not to say that all information should be targeted, but that more effort can be made in legislative drafting to bridge the disconnect between what 'looks like' a regulatory obligation, and what looks like a business feature.

Recommendation: Information provisions in consumer law should be subject to design obligations to communicate them effectively and dynamically, including within optimisation systems. This should lessen the divide between static regulatory requirements and the practical methods of information provision to users used by firms.

3.3.3 Functionality

Summary: The functionality layer refers to the more abstract capabilities of computing systems than interfaces, in this layer we find tasks that computing systems are designed to achieve for users, providers and others.

The functionality layer is deeper than the interface layer, although often is intertwined with it. Regulation aiming at the functionality layer obliges the creation of new or different underlying technologies, rather than just surfacing information about existing technologies, or providing options around such technologies that do not necessitate foundational redesign

or adjustment. In practice, interfaces could be considered functionality, while technologies such as recommender systems contribute to the functioning of interfaces, however we split the two by considering interfaces as the basic methods to view information and change settings provided by underlying functionality.

Functionality is often entwined with law. There are many functionalities related obligations we do not cover in this work, but which are relevant to platform regulation more broadly. Platform regulation has effectively obliged firms to create geolocation functionality to try and separate users by jurisdiction to treat them differently. Emerging regulation is demanding that firms create age verification or age ‘assurance’ tools to attempt to avoid certain features or content being accessible to minors. Here we zoom in on community stewardship functions however, which are the functions which interface with Accountability Groups and other users in order to attempt to ‘open up’ some of the operational functions and steering of optimisation systems — with mixed success and prospects.

3.3.3.1 Community Stewardship

Some legal provisions have obliged the creation of new functionality on online platforms that affects the optimisation stack.

According to Article 22 DSA, platforms should prioritise alerts to illegal content by trusted flaggers, acting within their designated area of expertise, provided the notice and action has been declared through the mechanisms set-out in Article 16. The status of trusted flagger is designated by the Digital Services Coordinator of the Member State in which the applicant is established. The powers offered might remain superficial, however. This is in large part due to the DSA’s focus on content, rather than patterns or issues that concern the optimization stack in its entirety. For example, a news recommender can reproduce social stereotypes (e.g. only display negative information concerning certain demographic groups) even if those singular pieces of news information are not “illegal content” (e.g. hate speech) per se.

Article 40 DSA (data access and scrutiny) does offer a more elaborate and strengthened form of scrutiny to verify the overarching design obligations of operators (described below), as well as the social and democratic impact their choices, including those that comply with the law, might have. More specifically, Article 40 DSA creates new obligations on platforms to provide information to vetted researchers and involves a back-and-forth process to get to an agreed form of data to provide. In practice, it implies the creation of new functionality as platforms must develop automated systems and APIs for researchers to mitigate the cost of dealing with each request manually. Article 39 DSA (Additional online advertising transparency) obliges VLOPs or VLOSEs that provide online advertising services to provide an interface through which these can be searched and examined. These provisions interact with the **Accountability Groups**, including academia and civil society groups, that sit alongside the stack. More specifically, they provide entry points to those groups to hold the digital ecosystem accountable when optimization strategies generate externalities on individuals, societies and environments.⁷² Effective scrutability is however contingent on there being sufficient well-intended

⁷² Bogdan Kulynych and others, ‘POTs: Protective Optimization Technologies’, Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (ACM 2020).

target groups who have the resources and expertise to scrutinise complex systems.⁷³ Still, unclarity remains. What happens when more structural issues are flagged? And more fundamentally, how easy is it for groups that discover unwarranted dynamics to flag them in the first place. Furthermore, given the many different laws that apply to interfaces, it might be difficult to identify which regulator to contact in case systemic risks are insufficiently addressed or which regulator should take the lead.

Recommendation: Trusted Flagging organisations should be able to raise complaints about the functioning of optimisation and recommendation systems to firms, rather than just about content.

Recommendation: Trusted Flagging organisations should also be able to submit a fast-track ‘super complaint’ to a body representing all relevant regulators in a jurisdiction, who should periodically discuss and assign responsibility for these issues amongst their (sometimes overlapping) competences. These should include consumer regulators, media regulators, data protection authorities, equality bodies, and the Digital Services Coordinator (if different).

Recommendation: Operators should explain and make available information that enables third-parties to test and scrutinise optimisation goals and their impact on/relationship to public values and societal interests (a right to observability and access). Such information, including auditable documentation standards, should not only pertain to the choices that impact the recommendation as such, but also provide insight into the data production processes and normative assumptions underlying recommender systems, which could include the labour conditions under which systems were developed, trained and deployed, and the instructions provided to team members regarding those data production processes.

3.4 Engine

Summary: Engines are software systems designed to fulfil optimisation logics, drawing on data to provide functionality, interfaces and more. Provisions targeting the optimisation engine require analysis and alteration of its functioning. Such systems directly require firms to tailor and alter the underlying technologies at the heart of optimisation — the engines that determine how optimisation functions should be calculated.

3.4.1 *Alternative engines*

As indicated above, interface provisions require platforms to allow users to switch between different engines where they exist. Article 38 DSA (Recommender systems) complements this provision, as it requires VLOPs and VLOSEs using recommender systems to provide at least one option not based on profiling. This might, for example, be a system that returns content chronologically, potentially further based on explicit choices of topics or individuals to follow.

⁷³ Jakko Kemper and Daan Kolkman, ‘Transparent to Whom? No Algorithmic Accountability without a Critical Audience’ (2019) 22 *Information, Communication & Society* 2081.

Article 38 DSA is relevant for the realisation of the right to (individual) self-determination in the sense that it must give users a real choice between different recommendation logics – one profiled and a non-profiled one. This provision will potentially give users a choice to compare how the recommendations look with and without profiling, even if the freedom to do so is a very limited version of a potential right to self-determination. The limitation of Article 38 DSA is, yet again, obvious when applying the stack perspective. While the provision focuses on the technical, the engine level, it does not have a corresponding obligation on the level of input data (in the sense of an option not to have one's personal data collected or processed to profile the consumer). Concerning the latter, Article 28 DSA acts as the sole exception, limiting the profiling and personal data processing of minors.

Other regimes, such as the GDPR in general, and Article 5(2) DMA (prohibition to combine data across platforms) more specifically, may interact with this requirement, but as these provisions cover multiple regulators, it is difficult to imagine these being coherently enforced. The GDPR could be interpreted as that alternative recommenders *have* to exist (e.g., those which use fewer personal data), and this may interact with the interface provisions to ensure that such alternative engines are surfaced to all users. Again though, to piece together these provisions across regulators requires a level of cooperation and coordination we do not currently see in the EU.

There are also no corresponding obligations on the operations level (in the sense of considering a preference of consumers whether to be profiled in the first place). As a result, it is again entirely up to the discretion of the platform or search engine provider to continue profiling as long as consumers have the choice to receive recommendations that are not based on profiling. Seeing that the core business logic that informs the design of social media architectures is to sell access to individuals and groups, such as through targeted advertising shaped by profiling, it is unlikely that Article 38 DSA will do much to impact the design of these architectures more profoundly.⁷⁴

Most importantly, an option that switches off profiling says nothing about the quality of the option that replaces it. If the alternative to a profiling-based recommender is one that sorts content in alphabetical order, this would undermine the point of the law. Recommenders that attempt to produce sensible and useful results without profiling users are possible, but the incentives to do so are not there. Firms are already litigating the DSA claiming that the opt-out alone would render irreversible economic damage to them (although the General Court has refused such arguments), and we will likely see this continue to be a contentious area.⁷⁵ Firms should be obliged to make alternatives for existing profiling methods and to evaluate them and continually improve them.

Recommendation: End-users should have a right to demand alternative recommendation options, and where technically feasible, a right to have third-party (stand-alone) recommenders to offer alternative recommendations. These too however, should align with the set of standards promoted as part of this study.

⁷⁴ Julie E Cohen, *Between Truth and Power: The Legal Constructions of Informational Capitalism* (Oxford University Press 2019).

⁷⁵ Case T-367/23 R *Amazon Services Europe Sàrl* ECLI:EU:T:2023:589 (Order).

Recommendation: Where engines that do not use profiling are selected, this should also trigger related obligations, such as a right to object under data protection law, ensuring that the underlying profiling never occurs, rather than is simply not shown.

Recommendation: As it stands, recommender rules allow users of large platforms to opt-out of seeing profiling-based recommendations, but provide no assurance that firms will invest capacity into running and maintaining useful and desirable alternatives. Structures which incentivise the production of these alternatives and align them with societal interests rather than private profit are required in all areas of EU law touching on optimisation.

Other regimes have a direct implication on the engines themselves. Article 17 GDPR (Right to erasure ('right to be forgotten')), read in conjunction with the chain of case law beginning with *Google Spain*,⁷⁶ has a direct effect on the requirements of certain optimisation systems produced by search engines. It requires them to have the functionality to remove certain records from appearing in these results, if they have succeeded in an erasure request. Further case law places new obligations on these systems. In particular, the *GC and Others* case requires operators to be able to alter their recommender system such that, whenever they receive a request for a delisting relating to a criminal proceeding, and regardless of the request's success, 'the overall picture [the list of results] [given to] the internet user reflects the current legal position, which means in particular that links to web pages containing information on that point must appear in first place on the list.'⁷⁷ The European Parliament version of the draft AI Act includes recommender systems of Very Large Online Platforms as 'high-risk' AI systems, which means there are obligations placed directly upon the design of the recommender. Such obligations include considerations of accuracy, bias, data representativeness, cybersecurity, and the potential for human oversight by the users of these systems. Under the current approach, these provisions are proposed to be elaborated by private standard-setting bodies CEN and CENELEC.

Further requirements for engine design may be derived from obligations relating to the processing of 'special category data'. Optimisation systems often place individuals and content inside a common geometric space, where content can be structured and grouped, alongside individuals, across thousands or even millions of 'dimensions'. Both content and individuals move as their optimal positions become apparent. Though these spaces can be constructed in abstract terms in the sense that they have no human interpretable meaning, the information they convey can nonetheless reveal interests or personal information that is deemed sensitive, such as a person's sexual orientation or political beliefs. Moreover, this structuring of content might occur automatically, without direct human interference or direction of the platform. In this context, the Court of Justice has confirmed that special categories of data also include data that would indirectly disclose 'following an intellectual operation involving deduction or cross-referencing'.⁷⁸ This conclusion places an indirect pressure onto operators to develop engines which can avoid this dynamic occurring, although whether that is possible at all remains an open question. At the same time however, Article 9 GDPR enables the processing of special categories of data when based upon explicit consent, which could result

⁷⁶ Case C-131/12 *Google Spain* ECLI:EU:C:2014:317.

⁷⁷ Case C-136/17 *GC and Others* ECLI:EU:C:2019:773, para 78.

⁷⁸ Case C-184/20, *Vyriausioji tarnybinės etikos komisija* ECLI:EU:C:2022:601.

in an overreliance of operators to use consent as their legal basis. Whereas Article 26(3) DSA prohibits profiling based on special categories of data, this prohibition only covers the presentation of advertisements. Finally, the logic established by the CJEU in *Meta Platforms* might entail that consent cannot be relied on when no viable recommender alternatives are provided to the end-users, whereby the notion “viable alternative” moreover cannot be interpreted to include the abandonment of the platform altogether.⁷⁹

Recommendation: Engines should be developed to promote self-determination and self-development for all, and the efforts thereto made available. To that end: a) the design team should be diverse, and, b) meaningfully engage with citizens and affected communities during the design process, including subsequent evaluation and iterations of the design, c) auditable documentation should be available as to how these interactions informed the design in a demonstrable or scrutable manner, d) metrics should be tested for their impact on people’s self-development and societal implications, for example for media pluralism. Design processes should include periodic considerations for public values and value-sensitive design strategies.

3.5 Input Data

Summary: We separate layers of the stack which deal with primarily expressive content from content which is primarily descriptive due to the different interests they entail. This divide is naturally imperfect and the data itself overlaps. The boundaries are also becoming less clear as traditionally expressive content such as images and text are now becoming automatically processable in.

In this Section, we make a distinction between considerations that concern the content of input data, and those that concern the users and environments affected by content data. There are some overarching recommendations that we believe could generally apply to the governance of input data:

Recommendation: An optimisation system producing diverse and varied outputs cannot do so without sufficiently diverse and varied inputs. Encouraging this should be the subject of multiple areas of regulation which happen in concert with changes to engines and interfaces to support a full-stack approach to governance.

Recommendation: Operators should provide end-users, research institutions and civil society, the ability to contest data and moderation decisions that could have a negative effect on the lawful representation of particular cultural or political groups, languages and forms of expression.

Recommendation: Operators should provide transparency on the content curation pool and the metadata they use.

⁷⁹ C-319/20 *Meta Platforms Ireland* ECLI:EU:C:2022:322.

Recommendations: In case operators make use of curation, they should consider the diversity of the content pool, and the impact current content pools have on marginalised and vulnerable individuals and social groups.

Recommendation: End-users should be given a functionality to curate profiles, including what personal information of theirs can be used, and give end-users the functionality to indicate, and make changes to, their own preferences and interests. When this is the case, end-users should be periodically reminded of the preferences and interests they chose and have the possibility to be redirected to an area where they can make changes to those preferences and interests.

3.5.1 Input Data (Content)

3.5.1.1 Limits of illegal content focus

Provisions concerning content have typically focussed on the removal of illegal content. Until the DSA, across Europe there were few obligations to remove content. Article 9 DSA constitutes a remarkable change in this regard, imposing orders to act against illegal content across Europe. Furthermore, Article 14(4) DSA obliges providers to act proportionately and in a way that respects fundamental rights when they enforce their own terms and conditions. This type of provision is emerging across the world, where laws operate via a two-step approach of a) mandating firms to place certain provisions in their terms and conditions, and b) obliging them to apply such terms consistently and proportionately.⁸⁰ This allows an expansion of the obligations in the DSA through the mechanism of obliging an alteration of terms of service, without the creation of new enforcement or monitoring structures. This regulatory tool however also re-enforces the privatisation of consumer regulation.

Some types of illegal trading are more difficult for platforms to detect or be reactive to than others. In addition, certain providers might have poor records of adhering to consumer rights, or uphold illegal practices around data collection, service and warranties, consumer information and similar. Such information may be held by consumer protection authorities, if they take action against these firms, or other bodies which retain records of compliance. Consumer bodies regularly provide tools and data ranking services, and third parties also aim to provide such data on trustworthy providers of goods and services. It may be worth allowing these factors to influence optimisation systems on large marketplaces.

Recommendation: Large commercial platforms selling goods and services through optimisation systems should be obliged to consider datasets produced by consumer organisations in rankings, as long as they are made to high standards. This should be initially explored through a code of practice.

⁸⁰ See e.g. Online Safety Bill (HL Bill 164 (as amended on Report), 2023, United Kingdom).

3.5.1.2 Data collection or labelling mandates?

Optimisation systems cannot produce diverse results without a diverse set of content, reflecting different perspectives, topics and genres. If a recommendation system were to contribute to the self-development and self-determination of a diverse set of users, it would need to be able to draw on a pool of information that is reflective of the diversity of its users. The approach that the DSA takes is a different one. As mentioned above, there are no quality requirements at the input level as long as the content is lawful. For example, there are no requirements that the pool of information must not be biased towards certain political views, cater to minorities or reflect different languages. Such requirements may be present in the AIA if the Parliament version passes, at least to some extent, as that regulation would place obligations around bias and the provenance of training data — although there is significant leeway for providers to interpret them as they see fit.

If more or less of certain content is desired — content of a certain ‘quality’; linguistic or cultural relevance; a political nature, or similar — then that content has to be labelled as such at some point in the stack. Doing so might be cheap or costly depending on the characteristics. Firms may not always wish to label content if doing so may bring obligations to change its distribution in ways that may not be profitable. For example, firms may not want to deter popular influencers from their platform by attempting to identify and label sponsored content that has not been declared. Similarly, when such labelling occurs, it is important that it is public and contestable. Labelling systems create ontologies of the world which can reify and reinforce biases.⁸¹

As it stands, there is no obligation in the DSA or similar regulations to actively label data such that other functions can be fulfilled. Labelling of data for downranking, shadow banning, reporting or on the basis of reports happens constantly. Running optimisation systems with certain preferences require and rely on such labelled data. Accountability groups can use such labels to study and scrutinise the functioning of systems.

Labelling mandates are complex to apply because they can be labour intensive, although systems like language models and computer vision claim to be able to reduce the cost of these. Such systems may be useful, but will need constant scrutiny of their performance and taxonomy due to the inherent challenges of algorithmic content moderation.⁸² A first step would be to publish these taxonomies and the labels on visible content.⁸³ Plenty of technologies already exist for this, many created by the industry itself, such as schema.org, which underpins the metadata consumers use for seeing e.g. opening times in services like Google or Apple Maps.

⁸¹ Geoffrey C Bowker and Susan Leigh Star, *Sorting Things Out: Classification and Its Consequences* (MIT Press 1999); Kate Crawford and Trevor Paglen, ‘Excavating AI’ <<https://excavating.ai>>; Abeba Birhane, Vinay Uday Prabhu and Emmanuel Kahembwe, ‘Multimodal Datasets: Misogyny, Pornography, and Malignant Stereotypes’ (arXiv, 5 October 2021) <<http://arxiv.org/abs/2110.01963>>.

⁸² Robert Gorwa, Reuben Binns and Christian Katzenbach, ‘Algorithmic Content Moderation: Technical and Political Challenges in the Automation of Platform Governance’ (2020) 7 *Big Data & Society* 2053951719897945.

⁸³ Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (December 2021), 86–92. <https://doi-org.kuleuven.e-bronnen.be/10.1145/3458723>; Milagros Miceli, Tianling Yang, Adriana Alvarado Garcia, Julian Posada, Sonja Mei Wang, Marc Pohl, and Alex Hanna. 2022. Documenting Data Production Processes: A Participatory Approach for Data Work. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 510 (November 2022), 34 pages. <https://doi-org.kuleuven.e-bronnen.be/10.1145/3555623>

Finally, where content is user generated, imposing direct requirements on platforms is hardly feasible, and platforms cannot or should not be expected to curate input content actively. However, this highlights the importance of the contestability of moderation decisions and the Terms of Use and Usage Policies those decisions are based on. The ability of, e.g. minority groups to effectively contest moderation or curation actions that unjustifiably limit the visibility of their content also matters for a right to meaningful optimisation.

Recommendation: Optimisation systems rely on explicitly or implicitly labelled data in order to function in line with user demands, and to meet legal requirements. Such labels should be publicly associated with records so they can be scrutinised. Taxonomies should also be publicly available to researchers and updated when they change in order so that the consequences of labelling the world in certain ways can be examined.

Recommendation: Whether the labelling of data is done internally or outsourced to third parties, it should be ensured that data work is performed under appropriate data labour standards as to avoid people's undue and unprotected exposure to mental and physical harm.

Recommendation: Content moderation decisions as well as the underlying Terms of Use themselves must be contestable and subject to scrutiny, including their proportionality and consequences for the ability of users to benefit from and exercise their fundamental rights.

3.5.2 Input data (users and environments)

User and environment data is a broad term that we use here to refer to data which does not relate to content. The term 'personal data' is inadequate for this because content may also be personal data of the individuals that it identifies and relates to. User data, in contrast, refers to the activities of users that are captured by actors within the optimisation stack. This may be data on clicks, views, follows, 'likes', telemetry data from sensors of devices like gyroscopes or Bluetooth, or data from the usage of other services that are captured by embedded trackers, such as 'pixel' trackers or software development kits (SDKs).

3.5.2.1 Sensitive data

As indicated above, individuals inside optimisation systems may be implicitly (and easily) grouped in ways which correlate to sensitive data categories. Article 9 GDPR (Processing of special categories of personal data) in effect requires private commercial entities to get explicit consent before the processing of certain categories of data that have been noted as relevant for recommender systems, such as sexuality, ethnicity or political opinion. While there are alternatives to consent, few will apply, and none routinely (see in this context, also our discussion above, including the pitfalls of relying on consent as a protective mechanism for individuals). Assessing the extent to which optimization systems create sensitive data from input data which has latent potential to be sensitive is a difficult task when labels for these

groups are not immediately available.⁸⁴ Some proposals for approaches which are privacy preserving exist, such as those using encrypted computation,⁸⁵ or those focussing on demographic data.⁸⁶ However, these methods need further exploration in context. The first step is to create knowledge of sensitivity in the context of optimisation systems, focussing on input data that may be permissible or impermissible to hold or collect.

The special categories of data show a strong connection with the “protected grounds” found in European non-discrimination law (the EU Equality Directives prohibit discrimination based on so-called race and ethnicity, sex, religion or belief, disability, age and sexual orientation). Yet, the efficacy of both applying and enforcing non-discrimination law in the digital environment has been questioned by scholars.⁸⁷ Among others, in discrimination law, a distinction is commonly drawn between direct and indirect discrimination on the basis of whether factors are explicitly used, or whether discriminatory effects can be observed. Unless expressly mandated, direct forms of discrimination are prohibited. Indirectly discriminatory rules or provisions can be objectively justified when they pursue a legitimate aim and the means of achieving that aim are appropriate and necessary. Due to the emergence of machine learning technologies, the conceptual distinction between direct and indirect discrimination has become increasingly blurred however: for instance, should the reliance on data that strongly correlates with a protected ground (i.e., proxy discrimination) be categorised as direct or indirect discrimination, and consequently, as justifiable or not? Moreover, in an era where people’s self-determination and self-development can be undermined based on a variety of characteristics, one can question whether a narrow focus on a select number of exhaustively enumerated grounds can be maintained. Furthermore, each equality directive has a narrowly constructed scope, both on a personal (the protected grounds they cover) and material (the (market) domains in which they apply) scope, further limiting the amount of contexts and settings in which recommender systems could be captured by them.

Recommendation: Digital Services Coordinators should proactively work with data protection supervisory, equality bodies and vetted researchers to identify data that is likely to be sensitive or discriminatory in the context of an optimisation system, and to either mitigate its sensitivity and abusive or discriminatory use at the point of collection or to restrict its use. Such an evidence base will support the enforcement of both regimes.

-
- ⁸⁴ Michael Veale and Reuben Binns, ‘Fairer Machine Learning in the Real World: Mitigating Discrimination without Collecting Sensitive Data’ (2017) 4 *Big Data & Society* 205395171774353.
- ⁸⁵ Niki Kilbertus and others, ‘Blind Justice: Fairness with Encrypted Sensitive Attributes’ in Jennifer Dy and Andreas Krause (eds), *Proceedings of the 35th International Conference on Machine Learning*, vol 80 (PMLR 2018) <<http://proceedings.mlr.press/v80/kilbertus18a.html>>.
- ⁸⁶ McKane Andrus and others, ‘What We Can’t Measure, We Can’t Understand: Challenges to Demographic Data Procurement in the Pursuit of Fairness’, *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (ACM 2021) <<https://dl.acm.org/doi/10.1145/3442188.3445888>>.
- ⁸⁷ Frederik J Zuiderveen Borgesius, *Discrimination, Artificial Intelligence, and Algorithmic Decision-Making* (Directorate General of Democracy, Council of Europe 2018). Jeremias Adams-Prassl, Reuben Binns and Aislinn Kelly-Lyth, ‘Directly Discriminatory Algorithms’ (2023) 86 *The Modern Law Review* 144. Janneke Gerards and Raphaela Xenidis, *Algorithmic Discrimination in Europe: Challenges and Opportunities for Gender Equality and Non Discrimination Law* (Publications Office of the European Union 2021) <<https://data.europa.eu/doi/10.2838/544956>> accessed 12 August 2021. Naudts (2023), *Fair or Unfair Differentiation? Reconsidering the Concept of Equality for the Regulation of Algorithmically Guided Decision-Making* (Doctoral Dissertation)

3.5.2.2 Device data

Some optimisation stacks have been known to seek access to data obtained by or stored on users' *terminal devices*, such as their smartphones. Such data might include information about how users act, or information about users' environments, such as from sensors on the device, fitness trackers, or similar. Article 5(3) ePrivacy Directive (hereafter ePD) on the confidentiality of communications renders access to such data prohibited if not necessary for a service requested by the user, or not accompanied by the explicit consent of the user. Consent under Articles 4 and 7 GDPR, to which the definition in the ePD is linked, requires it to be specific and separate from other matters. However, as it stands, data from users' devices seems frequently used within optimisation systems without attempts to gather such consent. The issue is compounded by the limited fines possible under the ePD, although some Member States, such as France, have heightened them in national law to match those possible under the GDPR. It is further compounded by dispersed competences — while some countries have data protection and e-Privacy law handled by the same regulator, others place it in the hands of another regulator, typically a telecommunications regulator.

As it stands, little progress is being made on an ePrivacy Regulation, which would ideally place the responsibilities for enforcement in the same location as the data protection authority. Therefore it is important that regulatory co-operation includes the ePrivacy regulators in each member state.

As with other forms of data, the legal obligation to obtain free and informed consent in terms of device data echoes forward and seemingly obliges firms to be able to run recommendation *without* such data. Digital Services Coordinators and the Commission should both ensure that recommenders without device data are available as the selectable 'options' in terms of the regulation at the interface layer.

Recommendation: Regulatory co-operation should include national e-Privacy regulators.

Recommendation: As it should be optional to provide device data to a service for the unnecessary purpose of recommendation, such cooperation should ensure that a recommender without device data is selectable by all users explicitly.

3.5.2.3 Data access

Access to data is a fundamental right in the European Union under Article 8 Charter, and is enshrined, inter alia, in the GDPR, the Digital Markets Act (as portability), and potentially in the forthcoming Data Act (in relation to specific devices).⁸⁸ Data access in some of these instruments is not limited to data that is observed about an individual, but also that which is inferred

⁸⁸ GDPR, art 15; Directive (EU) 2016/680 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and on the free movement of such data, and repealing Council Framework Decision 2008/977/JHA, OJ L 119/89 (Law Enforcement Directive), art 14; Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act) OJ L 265/1, art 6(9); European Commission, 'Proposal for a Regulation of the European

about them, such as opinions, whether from a computer or a human.⁸⁹ In theory, this would seem to allow an individual (or collectives of individuals) to inquire about what a recommender system ‘thought’ about them.

However, the data people get access to is often highly abstract in nature, as it has been turned into an ‘embedding’ — effectively a way of turning diverse sources of data into a way that can be used to compare people geometrically to how close they are to certain pieces of content, and each other. These embeddings have no direct human interpretation — they may be a vector of thousands of numbers between -1 and 1 that effectively represent coordinates in a high-dimensional space. On some of these dimensions, content may be close together (for example representing the fact they are in the same language, or in the same genre) while on other dimensions they may be further apart, in the various ways they are distinct. But this data is effectively *relative*. Having a copy of the embedding that describes a user or a piece of content is highly meaningful in the context of the wider infrastructure, but totally meaningless on its own, *because users have nothing to compare it to*. As a consequence, the right to access fails at its one of its main functions of allowing scrutiny and promoting accountability. This means that individuals looking for their digital ‘alter ego’ are unlikely to find it sufficiently using this approach.

Other approaches can be envisaged. Interactive services and explanations may be better at explaining the position of users within the data that is collected about them, but they would need proactive engineering to create. Regulators from multiple domains should collaborate to ensure these rights make sense within optimisation environments where data does not have human interpretable meaning.

Recommendation: Where services cannot provide effective scrutiny in response to rights of access due to the nature of the datasets and data systems, they should be obliged to invest in the creation of practical, interactive systems that serve an equivalent function, with suitable fidelity and detail.

3.6 Business-to-Business Interface

Summary: the ways that businesses interact with a service and in turn with consumers, which facilitate certain interactions and business practices over others. B2B interfaces are important as they mediate what information, settings and similar traders can provide in their interactions with platforms and online marketplaces. They have implications for liability across the supply chain as well as the location of the responsibility for compliance. They are important parts of the broader picture, and are regulated by the Fairness in B2B Platforms Regulation, and parts of the DSA, among others. We draw attention to it here but do not analyse it in detail, as it sits outside of the main focus of consumer groups.

Parliament and of the Council on Harmonised Rules on Fair Access to and Use of Data (Data Act), COM(2022) 68 final’ (23 February 2022), art 5.

⁸⁹ Jef Ausloos and Michael Veale, ‘Researching with Data Rights’ [2020] Technology and Regulation 136.

3.7 Connectivity Infrastructure

In certain systems, the underlying connectivity infrastructure can be important. For example, the governance of the optimisation systems of Facebook has been significantly changed by their [Internet.org/Facebook](#) Free Basics programmes, which increase reliance on those services and optimisation ecosystems compared to others, such as search engines on the open Web. These programmes are tie-ups with national telcos for zero-rating agreements for certain services, such as Facebook, while others outside the ecosystem, such as other apps and services, and access to search engines and their results more generally, cost more or require connection to WiFi.⁹⁰ In some jurisdictions, people reportedly use the Facebook search bar as a general search engine, which effectively heightens the reliance on certain types of optimisation as well as widens the number of use cases and societal role for these systems, heightening the governance stakes and scope.⁹¹ In the EU, zero-rating is prohibited by the Open Internet Regulation, the EU's net neutrality instrument.⁹² However, looking ahead, there is significant lobbying to undo and alter net neutrality in the EU, and the issue of connectivity infrastructure within the optimisation stack may become more relevant. This may particularly be the case in high mobile bandwidth areas such as virtual reality.

3.8 Whole-Stack Governance

Some regulatory provisions as well as some actors span the stack to a greater or lesser degree. These include

- **Operations and Management**
- **Accountability Groups, including civil society, (academic) research institutions, regulators and standardisation bodies**
- **Advertisers**
- **Individuals and communities.**

3.8.1 Operations and Management

Operations and Management encompasses the organisational aspects of platforms. They span the stack insofar as an actor is influential across it. A more centralised, vertically integrated platform might build the hardware, most of the apps, interfaces, collect the data, control the cloud services, and even run connectivity. Amazon, for example, builds a smart speaker (Echo), determines a lot of its interface (and steers the apps, or 'Skills', it does not), collects significant audio and transcription data, builds functionality with other 'smart home' tools in e.g. the Amazon Ring range, runs a huge proportion of the world's cloud services and machine learning tools, provides a business-to-consumer marketplace in the form of various online stores, and even runs its own telecoms network to link these devices together (Amazon Sidewalk). Other services, like Twitter/X, sit as an app with very little influence in domains of hardware or connectivity.

⁹⁰ Toussaint Nothias, 'Access Granted: Facebook's Free Basics in Africa' (2020) 42 *Media, Culture & Society* 329.

⁹¹ Peter Cihon and Helani Galpaya, 'Navigating the Walled Garden: Free and Subsidized Data Use in Myanmar' (*LIRNEasia*, 2017) <<https://perma.cc/22ED-CKKZ>>.

⁹² See e.g. Case C-5/20 *Vodafone (Tethering)* ECLI:EU:C:2021:676.

In practice, the operations layer is a layer of human problem-solving capacity to deal with often unexpected, undesirable and unintended effects of running optimisation systems. Recommender systems require constant maintenance and can't just be left running. To consider aspects such as fairness takes continuous analytic capacity and significant organisational effort.⁹³

Article 41 DSA (Compliance function) directly regulates compliance aspects of this layer, requiring compliance units of VLOPs and VLOSEs to be 'independent from their operational functions' and have 'sufficient authority, stature and resources [...] to monitor the compliance of that provider' with the DSA. The head of this function can only be removed by the management board of the VLOP or VLOSE. The GDPR or the draft AIA have no comparable direct requirements on firms, and it is notable these provisions apply only to VLOPs and VLOSEs — smaller entities do not have capacity obligations for their compliance functions.

This layer is ultimately responsible for compliance with individual aspects of governance in the stack, but it also has several overarching obligations which span these layers. Article 34 DSA (Risk assessment) requires VLOPs and VLOSEs to 'diligently identify, analyse and assess any systemic risks in the Union stemming from the design or functioning of their service and its related systems, including algorithmic systems, or from the use made of their services'. Article 37 DSA (Independent audit) requires the organisation of audits of certain DSA requirements, at those organisations' expense. Similarly, large-scale profiling of the sort carried out by entities in the optimisation stack is likely to trigger Article 35 GDPR (Data protection impact assessments), which require consideration of an extensive array of rights and freedoms. Other assessments may be required of users of recommenders as high-risk AI systems, if those provisions from the Parliament version of the draft AIA become law.

While, as we have seen, most consumer rights that the DSA offers concerning the regulation of recommenders do not go much deeper than the B2C interface level, the systemic risk provisions in Article 34 DSA are an example of a provision that potentially addresses the deeper layers of the recommender systems stack. The risk assessment must take place along the levels of the stack, including the B2C interface level (amplification and wide dissemination of illegal content), the operations level itself (decisions regarding the applicable terms and conditions), the input level (data-related practices of the provider), and the engine level (systems for selecting and presenting advertisements and design of the recommender system pursuant to Article 34(2)). Correspondingly, the mandatory mitigation measures can comprise different levels of the stack, including interaction with accountability groups such as trusted flaggers.

Articles 34 and 35 DSA acknowledge that recommendations can result from a complex interplay between different players and functionalities at different levels in the provision of recommendation services. Potential conflicts with, or failures to realise a right to self-determination and self-development can become relevant in the sense of systemic risks to the fundamental rights to human dignity, data protection, freedom of expression and non-discrimination to which Article 34 (1) explicitly refers. Having said so, identifying such a risk is left mainly to the platform's discretion. It does not include a corresponding obligation to design a recommender system in a way that promotes self-development and self-determination. This is likely a result of a difficulty in identifying a particular value to align with, but this can also be remediated by bringing representation, politics and accountability constructively *into* the process itself. The

⁹³ Henriette Cramer and others, 'Assessing and Addressing Algorithmic Bias in Practice' (2018) 25 *Interactions* 58.

operationalisation of complex and nuanced values cannot be done in one go, but requires a process — the politicisation of the algorithmic, and the opening up of future possibilities.⁹⁴

Recommendation: The DSA has a sufficient focus on risks but fewer focuses on positive, societally steered aims that optimisation systems might aim at. Participatory functions should be envisaged which create positive design obligations for actors behind optimisation stacks, rather than just a focus on avoiding downsides and pitfalls.

Recommendation: There should be organisational safeguards at the Operations/Management level to consider individual and societal interests, such as the right to self-determination and self-development, for example a dedicated role or team and processes to engage with external stakeholders and members of accountability groups as part of the process of defining optimisation goals. There should also be institutional support for individual and collective efforts to responsible recommender design, room for experimentation and learning and acknowledgement, for example as part of performance reviews.

In practice, the DSA seeks to provide additional capacity through Article 40 DSA (Data access and scrutiny) providing resources to external vetted researchers. However, the impact of this research — the closing of the loop from discovery to response or mitigation — is unclear. While Recital 90 DSA counsels that when carrying out risk assessments, providers should consider the state of the art, there is no obligation on providers to respond to the findings of researchers who have used their data in an open forum. Such a mechanism of dialogue is common in the security community, which has developed norms around responsible disclosure that rely on a dialogue with firms. A parallel obligation to foster such norms could be placed in this case.

Recommendation: Compliance and Risk Assessment Functions of organisations should have an obligation to consider and respond to research on their functionality. This obligation should require firms to publicly acknowledge issues discovered in optimisation systems, and publicly describe the actions they have and will take in response.

3.8.2 Accountability Groups

Where accountability groups are concerned, the EU's digital acquis concerning platforms relies heavily on external bodies for analysis and policing. Of particular relevance are the *vetted researchers* and *trusted flaggers* in Articles 22 and 40 DSA; the *representative bodies* in Article 80(2) GDPR; representatives of civil society and consumer protection organisations, and *worker groups* in Article 9 draft Platform Work Directive. Regulators too are accountability groups, and feature in most European legislation in this area in various forms and guises.

Such groups require capacity to operate. Funding is often scarce, and as such provisions that imagine a high-capacity sector just ready to go are often to face resourcing challenges

⁹⁴ Louise Amoore, *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others* (Duke University Press 2020).

undermining their effectiveness. This will be unevenly spread across the Union — certain jurisdictions simply have greater capacity in civil society than others.

Recommendation: Funding instruments for civil society and research organisations to hold optimisation systems to account should be considered. The independence of researchers that contribute to monitoring and enforcement actions must be respected and protected. The way research performances are evaluated and rewarded must be adjusted to acknowledge the activities that researchers engage in as part of their societal role under the new digital framework.

Other parts of the proposed digital *acquis* may provide inspiration for further provisions. The proposed Platform Work Directive creates roles for representatives to act on behalf of platform workers to scrutinise systems, with the support of technical experts if required.⁹⁵ Consumer bodies exist in every European jurisdiction and may play a similar role. In the draft Platform Work Directive, the ability to trigger an expert examination occurs when a significant change is made — this could be a trigger which requires e-commerce or similar platforms meeting a functionality and size threshold to consult with consumer bodies, and if the concern is significant, a cost transfer or cost sharing arrangement similar to the draft Platform Work Directive to fund an external expert with privileged access to systems could be envisaged.

Recommendation: The Commission should take inspiration from the draft Platform Work Directive and its provisions on supporting external expert analysis of platforms at a moment of significant change to consider a similar provision in relation to consumer bodies.

3.8.3 Individuals and Communities

Individuals and groups play a prominent role where rights are concerned. These may be existing rights, such as the right to access, the right to have certain recommender options, the rights to make complaints to platforms, or rights to judicial remedies against digital actors or their regulators.

To ensure whole-stack governance, regulators should establish an environment in which the actions and decisions of all actors that control the optimization stack, can be scrutinised. In order to establish a healthy digital ecosystem, processes should be envisaged that promote the democratic inclusion of those affected, or their representations.

Recommendation: During the ideation, design and development stage, operators should include meaningful consultation and representation of affected communities and consumers. Public consultations should be documented, including for what reasons citizens and communities were heard and how their feedback contributed to the design of the recommender systems in a demonstrable way.

⁹⁵ European Commission, 'Proposal for a Directive of the European Parliament and of the Council on Improving Working Conditions in Platform Work COM/2021/762 final' (9 December 2021), art 9(3).

Recommendation: Operators should enable third-parties (end-users, affected communities, civil society, etc.) to complain and make constructive suggestions.

Recommendation: Operators should explain and make available that enables third-parties to test and scrutinise optimisation goals and their impact on/relationship to public values and societal interests (a right to observability and access).

Recommendation: Operators should perform diversity, human rights, and systemic risk, impact assessment, taking into account individual, collective, social and democratic risks.

Recommendation: Operators should provide, either publicly, or through auditable documentation standards, insight into the data production processes and normative assumptions underlying recommender systems. This information could include the labour conditions under which systems were developed, trained and deployed the instructions provided to team members regarding the data production processes.

3.8.3.1 Right to be treated anonymously in commercial contexts?

Previous work from BEUC has indicated that it would be desirable to have a right to be treated anonymously in commercial contexts — or to ‘shop anonymously’. This would not mean that no information is collected about individuals (after all, they need goods paid for and shipped to them), but instead their commercial environment is not optimised. This is related to, but not identical to, ‘do-not-track’ proposals in data protection and e-Privacy law.

A few considerations need to be made in this respect. Firstly, being online, on a website or an app, is structurally not anonymous. Tracking mechanisms abound, many of them illegal in nature.⁹⁶ Insofar as these are integrated with recommendation and optimisation, the problem of shopping anonymously is entwined with the problem of *browsing anonymously*. As a consequence, it is key to pay attention to the recent judgment in *Meta Platforms*,⁹⁷ which tries to place some data protection firewalls between online tracking and experiences in platform optimisation systems. Enforcement of existing data protection law is necessary to make such a right to shop anonymously even possible in the first place.

A second relevant dynamic here is the supply chain of plugins, software development kits, APIs and similar that underpin much e-commerce. A constant challenge is that commercial sites may simply not be aware of the ways in which their ranking, recommendation and tracking systems function. This breaks some of the logics behind, for example, data protection law, which requires data protection by design in Article 25 GDPR, but does not extend this requirement to suppliers of data services and tools which European businesses integrate. As a result, it cannot be said that many good parts, or building blocks, can be combined to make a good — or even basically compliant — whole. Only by grappling with this supply chain, which requires thinking across the stack, can this right be realistically accommodated.

⁹⁶ Michael Veale and Frederik Zuiderveen Borgesius, “Adtech and Real-Time Bidding under European Data Protection Law” (2022) 23 German Law Journal 226.

⁹⁷ Case C-252/21 *Meta Platforms Ireland* ECLI:EU:C:2023:537.

Recommendation: A right to be treated anonymously in commercial context would require some foundational changes to the stack to be effective, including placing obligations on actors who are not data processor or controllers to only provide compliant or tracking-free tools.

While such a right to be treated anonymously could be accommodated quite easily for individuals who have logged in accounts, the logistics of this would need to be carefully considered for individuals who are browsing without logging in — presumably typical behaviour of individuals with these concerns, as consumers are already often advised to make use of ‘private windows’ and similar within Web browsers to avoid being judged and dynamically priced based on previous visits that have been associated with them through technologies such as cookies. One option is to present individuals with an option, like a pop-up or banner, upon entering every shop. This is likely to provoke significant backlash in an era of ‘consent fatigue’ (predominantly caused by firms attempting to continue to track illegally and gather invalid consent in the face of a *de facto* ban in European law). A better option, although an arguably more extreme one, is to forbid personalisation by default in situations where an individual is not ‘logged in’. This effectively would forbid the commercial use of third-party tracking where an organisation did not have a pre-existing relationship with a user. Businesses may state that this is a disadvantage to smaller organisations, as large online marketplaces would have a greater likelihood an individual would be logged in, and potentially consent to such optimisation. However, this issue can be remedied with a series of federated and interoperable logins, such as OAuth 2.0. These do not share data (e.g., tracking data), but instead allow authentication without constant registration. Individuals using these (or affiliating existing accounts to them) can durably set preferences for or against optimisation (or on preferences as to which types, they prefer) that follow them across contexts. Preferences such as this cannot be set in a browser without creating a ‘fingerprinting’ risk (i.e. allowing users to be more easily tracked against their will).

Recommendation: The best way to implement a right to shop anonymously is to prohibit tracking and optimisation of marketplaces where a user is not logged in — where that shop does not have a relationship with that user, and thus cannot ascertain durable preferences. To tackle the dominance related issues this might cause, the Commission should promote and invest in federated, interoperable logins, and require them to be supported in certain contexts.

Conclusion

Recommender systems are not simply a service or tool to push certain contents on users – recommender systems are the engines that enable and shape the digital experience and interactions of users and form an important part of the overall digital communications infrastructure. They have become an invaluable tool in our understanding and navigation of digital information society; a means to discover, learn and pursue the information goals people deem valuable economy. Seeing the central role of recommenders in the digital marketplace, we argue that users should not only have a legitimate interest to be protected from risks that recommenders pose to fundamental rights or other legitimate interests. Instead, recommender systems should be regulated with the goal of promoting self-development and self-determination *for all*; positively harnessed to dismantle, rather than reinforce (structural) inequalities,

or realise the self-serving interests of a select few private actors. End-users and non-end-users, as individuals or as members of social groups or collectives, too have a legitimate interest in optimisation strategies that are inherently useful and meaningful to them. The citizen-consumer should be accounted for, and be given (constructive) account within, and as part of, optimisation. Concretely this means a need to rethink the governance of the recommender system from the perspective of users and society. Citizen-consumers should:

- Be enabled to understand the rules of engagement, including how systems function, for which purposes they have been optimised, and the consequences such optimisation strategies entail on the content users see
- Have an actual and actionable say in the optimisation goals pursued within digital ecosystems
- Exercise meaningful choice and voice, which requires the presence of alternative options, both in relation to a particular recommender system's functioning and in relation to other operators, including service-providers and platforms
- Be included, represented and having one's voice heard and recognised during the ideation, design, deployment and evaluation of recommender systems, meaning they should have access to participation and contestation mechanisms.

To realise such a right to constructive optimization we also need to change the metaphors we use when talking about recommender systems. Instead of the popular 'black box' metaphor that has informed the way recommender systems are regulated in the DSA and the AI Act, we propose using the stack metaphor. Recommenders are not a box. They are the result of a complex dynamic interplay between different processes, technology layers and actors. Using the stack metaphor, we were able to demonstrate the futility of an easy regulatory quick fix to achieve constructive optimization but we also demonstrated that there are various intervention points, regulatory options and frameworks that can and should be used. The stack metaphor also highlighted the current fragmented nature of recommender governance, and the need for more consistency across the interpretation, application and enforcement of the different relevant frameworks, like the DSA, DMA, AI Act and the Platform Workers Directive.