

Primitives of Legal Protection in the Era of Data-Driven Platforms

Hildebrandt, Mireille

Published in:
Georgetown Law Technology Review

Publication date:
2018

Document Version:
Final published version

[Link to publication](#)

Citation for published version (APA):
Hildebrandt, M. (2018). Primitives of Legal Protection in the Era of Data-Driven Platforms. *Georgetown Law Technology Review*, 2(2), 252-273. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3140594

Copyright

No part of this publication may be reproduced or transmitted in any form, without the prior written permission of the author(s) or other rights holders to whom publication rights have been transferred, unless permitted by a license attached to the publication (a Creative Commons license or other), or unless exceptions to copyright law apply.

Take down policy

If you believe that this document infringes your copyright or other rights, please contact openaccess@vub.be, with details of the nature of the infringement. We will investigate the claim and if justified, we will take the appropriate steps.

PRIMITIVES OF LEGAL PROTECTION IN THE
ERA OF
DATA-DRIVEN PLATFORMS

Mireille Hildebrandt*

CITE AS: 2 GEO. L. TECH. REV. 252 (2018)

In this essay I discuss the political economy of data-driven platforms in terms of monopolies and monopsonies, arguing that the concentration of buying and selling power builds on and extends a pseudo-omniscient data architecture that feeds on an increasingly seamless data ecosystem. As the mathematical underpinnings of data-driven architectures are further extended into the hardware of smart homes, driverless cars and smart cities, they may at some point diminish or eradicate the semantic discontinuity that provides room for dissent and dissonance. Such discontinuity depends on a text-driven environment, whose affordances such as dissonance and adversariality cannot be taken for granted in the pervasive data ecosystem that drives the political economy of platforms. I contend that we need a small set of *primitives of legal protection* to shoot holes in the semantic continuity that is generated by the backend systems of data-driven platforms; competition law, consumer law, and tort law cannot contribute to the effective contestability of algorithmic decision-making by themselves. They require effective and practical rights to data minimization and purpose limitation and an actionable right to not be subject to automated decisions without a contestable justification. Once these primitives of legal protection are operational both the pseudo-omniscient nature of data architectures and the seamless nature of data ecosystems will be ruptured, thus reinstating the interstices for innovation, and for meaningful individual and political self-determination.

* Tenured Research Professor at the Faculty of Law and Criminology, Vrije Universiteit, Brussel. Part-time, full professor at the Science Faculty, Radboud University Nijmegen. I wish to thank Julie Cohen and Paul Ohm for inviting my contribution to the February 23rd, 2018 Symposium on *The Governance and Regulation of Information Platforms* and the other participants for their salient provocations.

I. THE POLITICAL ECONOMY OF DATA-DRIVEN PLATFORMS

The concern regarding platforms centers around their potential for monopolistic or totalitarian behavior. The lack of viable competitors not only disturbs the supposedly beneficial operations of a free market, it also exposes the users of such platforms to potentially monopolistic governance, leading them to accept terms of service across a number of services which leave these users vulnerable to unwarranted exposure and manipulation. Even if a platform does not intend to use its quasi-sovereign powers to actually institute a quasi-totalitarian rule across the many contexts we navigate, we should be concerned about its potential to do so. The difference between enlightened sovereign absolutism in 17th and 18th Century Europe and the institution of the Rule of Law in subsequent centuries was not that enlightened sovereigns did not mean well, but that we decided to subject their power to countervailing powers, while instituting democratic self-government. We, the people, want to decide for ourselves what we consider good or bad government. Such self-determination becomes even more critical as data-driven platforms have the potential to manipulate individual preferences and disrupt public opinion, underlining the need for a radical reconfiguration of the current data ecosystem. Though the ‘revelations’ about microtargeting by firms such as Cambridge Analytica, claiming to use “data to change audience behavior”¹ should not surprise anybody,² the uproar also nicely demonstrates the acuity of human audiences when confronted with the extent to which they may have been “played.”³ The time has come to redress the market power of oversized platforms (dinosaurs suffering from

¹ CAMBRIDGE ANALYTICA, <https://cambridgeanalytica.org> [<https://perma.cc/QL4B-N5F7>]. See, e.g., *Cambridge Analytica | UK news*, GUARDIAN, <http://www.theguardian.com/uk-news/cambridge-analytica> [<https://perma.cc/XP4X-FH95>].

² As extensively explored in M. HILDEBRANDT & SERGE GUTWIRTH (EDS.), *PROFILING THE EUROPEAN CITIZEN, CROSS-DISCIPLINARY PERSPECTIVES* (2008).

³ It may be interesting to compare the outrageously unscientific claims made for the effectiveness of marketing and advertising based on Freudian psycho-analysis, as described in 1957 by Vance Packard in *THE HIDDEN PERSUADERS*, with current claims about the effectiveness of micro targeting. See, e.g., ELI PARISER, *THE FILTER BUBBLE: WHAT THE INTERNET IS HIDING FOR YOU* (2011). Though both methods may be effective in the short term, I believe that human interaction is far too complex to predict, notably due to the anticipation of potential predictions that human interaction incorporates (coined ‘double contingency’). See, e.g., Raf Vanderstraeten, *Parsons, Luhmann and the Theorem of Double Contingency*, 2 J. CLASSICAL SOCIOLOGY, no. 1, 77–92 (2007). On the limited reach of machine learning, see, e.g., Rodney Brooks, *Machine Learning Explained*, RODNEY BROOKS: BLOG (Aug. 28, 2017), <http://rodneybrooks.com/forai-machine-learning-explained/> [<https://perma.cc/L67C-QBLL>] (notably the final paragraph).

data obesity and pattern obesity), while addressing the root of the problem at the level of data-driven societal infrastructures.

A. What Is a Platform?

This essay refers to platforms whose functionality, power, and affordances depend on the data-driven architecture they embody, and on the data-driven ecosystem they enable and configure. Data-driven platforms are neither mere corporations nor states, nor are they mere technical scaffolding.⁴ They afford both markets and political discourse, as with the Greek *agora* and the Roman forum. Both the *agora* and the forum originated as open spaces used for both commercial negotiation and political discourse. The Roman forum, however, achieved a certain elevation (it was in fact built on a platform) and turned into a more ceremonial, formal and enclosed terrain during the reign of the empire.⁵ Whereas the *agora* underlines the equality of the Greek freemen, the *forum* became the site for a political architecture that separated the emperor and his entourage from his subjects.

As Gillespie notes, the term platform was originally used in the computational sense of “a technical base upon which other programs will run,” while at some point, “[p]latforms are platforms not necessarily because they allow code to be written or run, but because they afford an opportunity to communicate, interact or sell.”⁶ As with the *agora* and the forum, the physical underpinnings of the site afford a mix of commercial transactions, artistic expression, and political discourse which highlights the *role* of data-driven platforms rather than their *computational architecture*. It is, however, the data-driven architecture and the data ecosystem it enables that determine what roles are possible and how they are distributed and constrained. This also relates to the question of whether these platforms are intermediaries (merely enabling content provision), media (responsible for the content itself) or—as Gillespie proposes—moderators⁷ (tasked with safeguarding what I would call the civility or even “civiness” of online discourse, raising difficult questions about

⁴ Tarleton Gillespie, *The Politics of ‘Platforms,’* 12 NEW MEDIA & SOC’Y 349–55, 355–57, 358–59 (2010) (349–55 on the myriad meanings of ‘platform’, 355–57 on the policy aspects, and 358–59 on the role of platforms as curators and mediators of public discourse).

⁵ Félix Ortega, *The New Public Space of Politics*, 14 INT’L REV. SOC. 209 (2004).

⁶ Gillespie, *supra* note 4, at 351.

⁷ See Tarleton Gillespie, *Regulation of and by Platforms*, in THE SAGE HANDBOOK OF SOC. MEDIA 257–58 (Jean Burgess et al. eds., 2017), <http://sk.sagepub.com/reference/the-sage-handbook-of-social-media/i2081.xml> [<https://perma.cc/9Q3Q-XXEJ>].

censure and the privatization of law enforcement).⁸ However, the distinction between these three modi of platform behavior is becoming increasingly artificial, as platforms incorporate many other behaviors, functions, and roles, notably providing many of the services they must supposedly moderate.⁹ This suggests that the framing of the question—whether platforms are intermediaries, media or moderators—is mistaken and distracts attention from what should concern us; the potential for exposure and manipulation that is inherent in data-driven platforms.

Due to their distributed, networked, and data-driven architecture, platforms enable the construction of invasive, over-complete, statistically inferred, profiles of individuals (exposure), the spreading of fake content and fake accounts, the intervention of botfarms and malware as well as persistent AB testing, targeted advertising, and automated, targeted recycling of fake content (manipulation). In this essay, I refer to platforms as (1) layered technological architectures that enable data-driven applications based on various types of machine learning, and (2) as corporations that incorporate vertical and horizontal competitors into their own legal architecture. First, I explain the interplay between both types of platforms and then rearticulate the issue of platforms in terms of the semantic continuity they generate as the biggest threat to privacy as identity construction,¹⁰ and to our capability to access, assess, redistribute, discuss, and rearticulate meaningful information. Finally, I propose a small set of primitives of legal protection as gatekeepers for a robust data-driven architecture that builds on and sustains a discontinuous polyphonic data ecosystem that allows for dissonance, dissent, and reinvention.

J. Monopolistic and Monopsistic Behaviors

By now, the term platform is often used to depict a new type of corporation that spreads its wings both horizontally and vertically to incorporate ever more data-driven services. This holds for the big five, GAFAM (Google, Apple, Facebook, Amazon, Microsoft), but also for Twitter, for corporations involved in the sharing economy (Uber,

⁸ Tarleton Gillespie, *Platforms are Not Intermediaries*, 2 GEO. L. TECH. REV. 198 (2018).

⁹ The antitrust proceeding of the European Commission against Google is based on allegations that its moderation ended in an unfair trade practice. See Summary of Commission Decision of 27 June 2017

Relating to a Proceeding Under Article 102 of the Treaty on the Functioning of the European Union and Article 54 of the EEA Agreement, 2018 O.J. (C 9) 11, 13–14 (fining Google €2,424,495,000 and prohibiting it from further infringements).

¹⁰ Privacy as “the freedom from unreasonable constraints on the construction of one’s identity.” Cf. PHILIP E. AGRE & MARC ROTENBERG, *TECH. AND PRIVACY: THE NEW LANDSCAPE* 7 (2001).

AirBnB), and for those traditionally focused on hardware such as IBM (now deep into machine learning with its Watson platform).¹¹ Platform corporations move from hardware to software, from mediation to service provision to moderation, and from software to hardware, notably in the form of cyberphysical infrastructures (Internet of things). Platformed corporations challenge the law's capabilities and legislators' willingness to sustain checks and balances in both global and national markets, noting, for instance, that current competition law fails to address the kind of issues raised.¹² This has sparked debates about a so-called platform economy,¹³ suggesting that the rise of data-driven platforms generates new types of affordances for citizens, users, consumers, customers, corporations, self-employed service providers, governments, legislators, regulators, publishers, advertisers, and others. Some proclaim the advent of a new type of capitalism, coining it 'platform capitalism,'¹⁴ building on earlier work concerning informational¹⁵ and surveillance¹⁶ capitalism.

In economic terms, market failure seems operative in three ways. First, consumers are paying an excessively high price for free services, because they pay for them with their personal data and ultimately with their privacy and manipulability. Instead of cheering about the low costs of services, Dunne notes that "the indirect cost of services within the digital economy—particularly the surrender of valuable personal data—comprises a form of excessive pricing" (emphasis omitted).¹⁷

¹¹ *It Knows Their Methods. Watson and Financial Regulation*, ECONOMIST (Oct. 22, 2016), <https://www.economist.com/finance-and-economics/2016/10/22/it-knows-their-methods> [<https://perma.cc/DPK4-S2UL>].

¹² Niamh Dunne, *Competition Law (and Its Limits) in the Sharing Economy*, in CAMBRIDGE HANDBOOK OF THE LAW OF THE SHARING ECONOMY (forthcoming 2018). However, data protection law may offer more adequate tools to deal with market failures due to the power of platforms, see e.g., *Facebook Loses Belgian Privacy Case, Faces Fine of Up to \$125 Million*, REUTERS (Feb. 16, 2018), <https://www.reuters.com/article/us-facebook-belgium/facebook-loses-belgian-privacy-case-faces-fine-up-to-125-million-idUSKCN1G01LG> [<https://perma.cc/S7FA-BG67>].

¹³ See generally Martin Kenney & John Zysman, *The Rise of the Platform Economy*, 32 ISSUES IN SCI. AND TECH. 32016, <http://issues.org/32-3/the-rise-of-the-platform-economy/> [<https://perma.cc/RA2N-JEGE>].

¹⁴ NICK SRNICEK, PLATFORM CAPITALISM (2016) (arguing that a new type of data-driven capital accumulation has transformed the nature of capitalism, resulting in new types of transnational power-imbalances).

¹⁵ See generally Julie E. Cohen, *The Regulatory State in the Information Age*, 17 THEORETICAL INQUIRIES IN L. 369 (2016), <http://www7.tau.ac.il/ojs/index.php/til/article/view/1425> [<https://perma.cc/2MVY-RK5D>].

¹⁶ See generally Shoshana Zuboff, *Big Other: Surveillance Capitalism and the Prospects of an Information Civilization*, 30 J. INFO. TECH. 75 (2015).

¹⁷ Dunne, *supra* note 12, at 9.

The indirect costs basically refer to the consequences of a business model that thrives on data-driven, behavioral advertising, geared towards nudging potential consumers into specific types of clickstream behaviors (thus generating revenue for the platform). As others have noted,¹⁸ the reconfiguration of online platforms into sites that are conducive to making people click on advertisements has simultaneously reconfigured them for sentiment analysis, opinion mining, and nudging as to political opinion.

Second, platforms that enable the sharing economy broker between consumers on the one hand and a plethora of small, independent service providers on the other, may lower the price of such services for consumers, but simultaneously falsely turn employees into independent providers without any bargaining power, hiding market failure on the labor market: “Conversely, the principal objection to sharing economy platforms for traditional providers—their propensity to undercut on price—is unlikely to find purchase as a theory of harm. While sustained below-cost pricing is illegal, taking advantage of a lower cost base emphatically is not.”¹⁹

Third, the labor market is suffering from monopsony, i.e., anti-competitive concentration in a labor market where a few firms dominate hiring in that market. Whereas a monopoly concerns concentration in the seller’s market (for instance, meaning that the sellers can set a high consumer price for the product or service), a monopsony concerns concentration in the buyer’s market (for instance, meaning the buyers can set a low price for labor or work). “Antitrust enforcement is mainly concerned with consumer welfare, and hence the impacts of a lack of competition on product prices, not wages.”²⁰

¹⁸ E.g., *The Turing Test Isn’t Won by Machines, It’s Lost by Humans*, SMERITY.COM (Jan. 8, 2018), http://smerity.com/articles/2018/turing_test_isnt_won_its_lost.html [<https://perma.cc/LQV9-3T64>] (“Any system or network that has been optimized for advertisements has been implicitly optimized for spreading misinformation.”); Zeynep Tufekci (@zeynep), TWITTER (Feb. 16, 2018, 10:54 AM), <https://twitter.com/zeynep/status/964573636530024448> [<https://perma.cc/8776-26CN>] (“Indictment shows RU used social media just like any other advertiser/influencer.’ They used the platforms as they were designed to be used to make money. I watched this up to the election. Such deliberate, false incitement was not just done by foreigners.”); see also *United States v. Internet Research Agency*, No. 1:18-cr-00032-DLF, 2018 WL 914777 (D.D.C. Feb. 16, 2018), <https://www.justice.gov/file/1035477/download> [<https://perma.cc/9DTR-AGC7>]; Zeynep Tufekci, *Engineering the Public: Big Data, Surveillance and Computational Politics*, 19 FIRST MONDAY (2014), <http://firstmonday.org/ojs/index.php/fm/article/view/4901> [<https://perma.cc/293F-BVX2>].

¹⁹ Dunne, *supra* note 12, at 10.

²⁰ José Azar et al., *Labor Market Concentration*, IZA – INST. OF LAB. ECON. 1 (2017).

The basic idea of monopsony is that if employers don't have to compete with one another for workers, they can pay less, and workers will be stuck without the outside job offers that would enable them to claim higher wages. The monopsony story is consistent with a wide range of observed labor market phenomena: wage stagnation, declining geographic and job-to-job mobility²¹

If market failure on the side of the labor market is not redressed, consumers will, in the end, be worse off for the simple reason that most consumers are also employees or small independent service providers. Though consumer price may be said to go down (if one abstracts from the "price" paid in personal data), most consumers will see their wages stagnate or diminish, unless the legal framework enables courts to intervene, as the Court of Justice of the European Union (CJEU) did in its ruling on Uber, though this was not based on competition law but on EU legislation concerning the provision of transport services.²²

²¹ Marshall Steinbaum, *How Widespread Is Labor Monopsony? Some New Results Suggest It's Pervasive*, ROOSEVELT INST. (Dec. 18, 2017), <http://rooseveltinstitute.org/how-widespread-labor-monopsony-some-new-results-suggest-its-pervasive/> [<https://perma.cc/UP3C-AA2C>].

²² *Asociación Profesional Élite Taxi v. Uber Systems*, No. C-434/15, 2017 E.C.R. 48, Opinion of Advocate General, Court of Justice of the European Union [CJEU], (May 11, 2017). (“[A]n intermediation service such as that at issue in the main proceedings, the purpose of which is to connect, by means of a smartphone application and for remuneration, non-professional drivers using their own vehicle with persons who wish to make urban journeys, must be regarded as being inherently linked to a transport service and, accordingly, must be classified as ‘a service in the field of transport’ within the meaning of Article 58(1) TFEU. Consequently, such a service must be excluded from the scope of Article 56 TFEU, Directive 2006/123 and Directive 2000/31.”) The CJEU C-434/15, 2017 E.C.R. 39, (Dec. 20, 2017) grounds its qualification of Uber’s service as a “transport service” instead of merely “an information society service” on the fact that

[Uber] exercises decisive influence over the conditions under which [that] service is provided by those drivers. On the latter point, it appears, inter alia, that Uber determines at least the maximum fare by means of the eponymous application, that the company receives that amount from the client before paying part of it to the non-professional driver of the vehicle, and that it exercises a certain control over the quality of the vehicles, the drivers and their conduct, which can, in some circumstances, result in their exclusion.

K. The Matter of Concern: A Political Economy of Machinic Neoplatonism

The political economy of platforms thrives on a semi-totalitarian *capture* of behavioral data,²³ while the seamless access to such data informs the economic power of platforms due to the knowledge asymmetry it entails. Due to their dependence on data-driven decision-making, platforms are rooted in a worldview that is deeply mathematical, whether implicitly or explicitly. As McQuillan notes, this entails a type of neoplatonism; seeking reality in the mathematical functions that supposedly underlie human behavior.²⁴ Mark Zuckerberg has revealed that one of his focal interests is in finding the “fundamental mathematical law underlying human social relationships that governs the balance of who and what we all care about.”²⁵ Though this may sound wildly off key to some, it actually refers to a key assumption of machine learning, which operates on the hypothesis that the mathematical functions by which datapoints are correlated are an approximation of mathematical laws that are believed to underlie patterns in real life. The latter belief forms the foundation of machine learning. Even if it were not true, it may become true if platforms base a plethora of machine-to-machine and machine-to-human decisions on the mathematical functions they elicit from the datapoints to which they have access. This is the machinic version of Merton’s famous Thomas theorem: if machines define a situation as real, it is real in its consequences.²⁶

Next to their online data-driven decision systems, most of the platforms are moving into the Internet of things, investing in cyberphysical infrastructures such as cloud robotics, connected cars, warehouse logistics, drones, smart homes, and smart cities. Insofar as such cyberphysical infrastructures are based on machine learning, they entail

²³ See Philip E. Agre, *Surveillance and Capture: Two Models of Privacy*, 10 THE INFO. SOC’Y 101, 118–22 (1994), where Agre saliently discusses the computer science term ‘capture’ as a means to understand how computer systems reconfigure human society in their need to obtain ever more data.; see also, Mireille Hildebrandt, *Privacy as Protection of the Incomputable Self: From Agnostic to Agonistic Machine Learning*, 20–1 THEORETICAL INQUIRIES IN L. 12–14 (2019 forthcoming), <https://papers.ssrn.com/abstract=3081776> [<https://perma.cc/829U-R7J6>].

²⁴ See Dan McQuillan, *Data Science as Machinic Neoplatonism*, PHIL. & TECH. 9 (2017).

²⁵ Michael Rundle, *Zuckerberg: Telepathy is the Future of Facebook*, WIRED UK (July 1, 2015), <http://www.wired.co.uk/article/facebook-zuckerberg-qa-the-future> [<https://perma.cc/342B-DMET>].

²⁶ See Robert K. Merton, *The Self-Fulfilling Prophecy*, 8 THE ANTIOCH REV. 193, 193 (1948); MIREILLE HILDEBRANDT, SMART TECHNOLOGIES AND THE END(S) OF LAW: NOVEL ENTANGLEMENTS OF LAW AND TECHNOLOGY 197 (2015).

what McQuillan has coined “machinic neoplatonism”²⁷ reconfiguring our “onlife”²⁸ world into an entirely different kind of society, structured upon radically different premises compared to those of text-driven institutions such as representative democracy and the Rule of Law.

My concern is whether, and if so, how, the data-driven architecture of platforms can afford individual human flourishing, autonomy, meaningful choice, agonistic democratic decision making, fair distribution of risk and opportunities, and balanced growth of public and private welfare. This requires keen attention to the relationship between their technical and corporate architecture on the one hand and what Julie Cohen has described as semantic discontinuity on the other hand,²⁹ taking into account the mathematical underpinnings of their operations. The issue with platforms is that they have become constitutive of human society in a number of ways and will continue to reconfigure and reconstitute employment, education, political discourse and decision making, healthcare and medicine, and the supply and distribution of energy.³⁰ The extent to which everyday life has become rooted in an *onlife* world run by a set of global platforms raises the question to what extent our agency is being diminished, overruled and tweaked in function of a corporate profit that is based in the quest for ever more behavioral data to further approximate the mathematical functions that supposedly determine our interactions.

Platforms redefine the space (1) for generating new meaning, and (2) for experimenting and playing around with new ways of seeing the same thing. They also redefine (3) the space for resistance against old and new ways of addressing us as parents, employees, students, employers, passengers, drivers, consumers, citizens, or service providers. As a result of combining a seamless data-driven pseudo-omniscience with a sequence of market failures, we may be facing the loss of semantic discontinuity³¹ that is key to a society that fosters and protects the capabilities of individual human beings, their participation in a viable democracy, and the checks and balances of an effective Rule of Law.

²⁷ See generally McQuillan, *supra* note 24.

²⁸ THE ONLIFE MANIFESTO: BEING HUMAN IN A HYPERCONNECTED ERA (Luciano Floridi, ed. 2014); HILDEBRANDT, *supra* note 26.

²⁹ JULIE E. COHEN, CONFIGURING THE NETWORKED SELF: LAW, CODE, AND THE PLAY OF EVERYDAY PRACTICE 239–41 (2012).

³⁰ This is not to say that they will supply energy or provide healthcare; I am referring to the way they are reconfiguring the way energy usage and other critical infrastructure is managed (e.g., via Google’s Nest).

³¹ COHEN, *supra* note 29.

II. SEMANTIC DISSONANCE IN THE ERA OF DATA-DRIVEN PLATFORMS

In her seminal work on the networked self, Julie Cohen discusses the protective potential of semantic discontinuity, demonstrating how it protects the practice (not merely the idea) of privacy, by offering interstitial spaces to reinvent the self.³² Such reinvention is not a matter of voluntary sovereignty; the self is not the assumed neoliberal subject that goes around reinventing herself in the process of exercising her consumer choice, or in the process of enhancing her competitive advantages. Reinvention happens in the gaps and frictions between what people mean and how others “read” them, keeping in mind that meaning is always about actionable mutual expectations.³³ It matters what we mean and how we are read, because this orients our interactions and those of others. Cohen highlights the playful character of self-construction that depends on the freedom to give meaning to the world, the self, and to the choices one makes or cannot make. In that sense, freedom does not depend on operationalizing given preferences, but on developing and reconfiguring the skills to “make a life.” It is about capabilities rather than mere opportunities. Taking Zizek to heart, we could say that “communication is a “successful misunderstanding.””³⁴ To actually *work*, the misunderstanding must be successful; it must set things in motion. Complete congruity between intent and response is not given, and attempting to freeze them into conformity would result in a stalemate and erode the generative nature of incongruent understandings. It would thus limit the “hypotheses space” of human understanding, diminish human imagination, and thereby reduce our capability to find new ways of dealing with real world problems. It would, in the end, also stifle innovation.

A. Text-Driven Dissonance

Just like physical walls and other tangible barriers, the mutual expectations we call social norms and the institutions they sustain³⁵ can offer a kind of “natural” protection against invasive inquiry. More importantly, institutions and social norms can provide resilience against the imposition of meaning that rules out any kind of play. They allow

³² *Id.*

³³ *Cf.* CHARLES SAUNDERS PEIRCE, *The Pragmatist Maxim and Wittgenstein's Understanding of Meaning*, in *SELECTED WRITINGS* (Philip P. Wiener, ed., 1958).

³⁴ SLAVOJ ZIZEK, *LOOKING AWRY: AN INTRODUCTION TO JACQUES LACAN THROUGH POPULAR CULTURE* 30 (1991).

³⁵ *See* PETER WINCH, *THE IDEA OF A SOCIAL SCIENCE* 60, 128 (Routledge & Kegan Paul 1958).

incremental as well as radical reinvention of the self. Though such material, social, and institutional resilience may be seen as “natural” as long as it is taken for granted (naturalized), the artificial character of institutions comes to the fore when less flexible computational constructs take over. Whereas our current institutions are largely contingent upon text-driven information and communication technologies (ICTs, namely the technologies of the word),³⁶ data-driven platforms may not share the affordances of technologies such as the printing press.³⁷ The latter induces a linear, sequential processing of information (i.e., from left to right, from top of the page to the bottom, from the first to the last page), while simultaneously forcing the reader to take note of a host of different viewpoints due to the proliferation of books written by authors with a variety of potentially contrarian perspectives. Text-driven ICTs, in point of fact, induce slow thinking (reflection), as well as a permanent anticipation of other ways of thinking the same things (imagination). The need for interpretation and the potential for adversarial behavior are baked into the affordances of written and printed text. Though we need not be naïve or romantic about the affordances of a text-driven society, both in-depth reflection and complex abstraction have been attributed to the externalization of memory and the mediation of reasoning that is enabled by text (including both letters and numbers).³⁸

In previous work I referred to the foundational poly-semantics of human language to explain the difference between machine translation and human translation.³⁹ Where machine translation thrives because the poly-semantics can be “resolved” and reduced to a univocal decoding and recoding, human translation thrives when the poly-semantics of an expression in one language is transposed into the poly-semantics of another language, while remaining loyal to the original expression, taking

³⁶ Compare WALTER ONG, *ORALITY AND LITERACY: THE TECHNOLOGIZING OF THE WORD* 77–136, 148–52 (Methuen 1982) (explaining how writing restructures consciousness as well as the culture it constitutes), with Jack Goody & Ian Watt, *The Consequences of Literacy*, 5 *COMP. STUD. IN SOC’Y AND HIST.* 332–44 (1963) (describing the multifarious impact of record-keeping, multiplication and categorisation of lines of argumentation and the increased complexity of societal order this entails).

³⁷ PIERRE LÉVY, *LES TECHNOLOGIES DE L’INTELLIGENCE: L’AVENIR DE LA PENSÉE À L’ÈRE INFORMATIQUE* (1990).

³⁸ JAMES GLEICK, *THE INFORMATION: A HISTORY, A THEORY, A FLOOD* (2010); ELISABETH EISENSTEIN, *THE PRINTING REVOLUTION IN EARLY MODERN EUROPE* (2005).

³⁹ Mireille Hildebrandt, *Law as Information in the Era of Data-Driven Agency*, 79 *THE MOD. L. REV.* 26–27 (2016); see also Douglas Hofstadter, *The Shallowness of Google Translate*, *ATLANTIC* (Jan. 30, 2018), <https://www.theatlantic.com/technology/archive/2018/01/the-shallowness-of-google-translate/551570/> [https://perma.cc/WME6-AWZW].

note that this can be done in multiple ways.⁴⁰ It is the notion of loyalty that is crucial here, and the awareness that such loyalty does not reduce to one necessarily best translation. It is such loyalty that distinguishes semantic dissonance from a false note. A translation can be developed in counterpoint to the original, but the potential for dissonance should not lead to disloyalty, which would imply that the original expression is taken for a ride in its alleged translation. Semantic dissonance is essential for a viable democracy, as it opens new ways of seeing the same thing, generating a larger repertoire for the framing of problems and the construction of solutions. As Ricoeur demonstrated in his work on the use of metaphor,⁴¹ a successful and sustainable society thrives on the new insights that are generated in the interstices of stabilized meaning, noting that a new metaphor allows one to see the same thing from a different perspective and may even lead to seeing new “things”. Whereas a false note is not tuned to other notes, the dissonant is; whereas the false note creates chaos, the dissonant creates tension. Dissonance thus adds two elements to the semantic discontinuity on which it depends. First, it is generative in its contrapuntal progression, enabling a coupling of consonance and dissonance that ties different melodies into a polyphonic movement. Second, it enables distinguishing productive (dissonance, polyphony) from destructive meaning generation (false notes, cacophony). This raises the pertinent question how much cacophony a society can support before breaking down into chaos, and when univocality will result in the weakness of a tyranny or totalitarian system that ignores opposition. Both questions are highly relevant for the reign of platforms. Do data-driven platforms reconfigure the balance between semantic continuity and discontinuity and the ensuing openings for dissonance and discord?⁴² Does the architecture of cross-platform tracking and inferencing result in an overall continuity that may not even be semantic, but remains stuck at the level of machinic behaviors? Will the combination of exposure and

⁴⁰ Hofstadter, *supra* note 39.

⁴¹ See generally PAUL RICOEUR, *THE RULE OF METAPHOR: THE CREATION OF MEANING IN LANGUAGE* (3d ed. Routledge 2003).

⁴² Note that the difference between a dissonant and a false note can only be established in the context of the tonality (the temporal progression) of a piece of music. See Ian Gallacher, *The Court's Dilemma: Or, Harmony and Dissonance in Legal Language*, 9 LEGAL COMM. & RHETORIC: JALWD 15 (2012), <http://www.alwd.org/lcr/archives/fall-2012/gallacher/> [<https://perma.cc/6TFA-TXUR>]. Other than Gallacher suggests, this makes the difference even more interesting and the metaphor of dissonance even more crucial for generating meaning. Musical ‘meaning’ emerges on the nexus of the vertical (contrapuntal) and the horizontal (tonal), and it is the semantic discontinuity that is inherent in human language and human institutions that enables contrapuntal and tonal switches in human communication (both polyphonic and cacophonic) to shift from one melody to another, e.g., turning a false note into a dissonant.

manipulability result in a loss of counterpoint, a lack of vital dissonance, and a monocultural equalization?

Semantic discontinuity affords dissonance. In that sense it is generative as it allows for shifts in meaning and for new meaning, regardless of intent.⁴³ At the same time it protects the self from overdetermination by the imposition of meaning by other humans or institutions; the affordances of text-driven institutional settings enable the self to develop other, potentially contrarian interpretations of words, phrases, institutional or legal norms different from those presented as obvious or final. The semantics of text are inherently poly-semantic; the ensuing contestability cannot be ruled out by force, even where acting on new interpretations may be prohibited or rendered impossible. It is not clear whether and how data-driven architectures have similar affordances. To the extent that the political economy of machinic neoplatonism drives platforms towards a web of contiguous data ecosystems which act on a buzzing stream of data-driven predictions, spaces for semantic dissonance may be seriously reduced or even closed.

A. Data-Driven Closure?

Though computing systems have unprecedented abilities to store bits and bytes and to calculate and correlate, their achievements are based on simulation rather than invention. Machine learning is not a matter of reasoning but of inductive calculation, of translating the relationship between input and output data into a mathematical target function that supposedly underlies the relationship between human experience and human judgment.⁴⁴ Basically, data-driven systems parasite on the expertise of domain experts to engage in what is essentially an imitation game.⁴⁵ There is nothing wrong with that, unless we wrongly assume that

⁴³ PAUL RICOEUR, INTERPRETATION THEORY (1976).

⁴⁴ THOMAS MITCHELL, MACHINE LEARNING 7–9 (1997).

⁴⁵ See A.M. Turing, *Computing Machinery and Intelligence*, 49 MIND 433 (1950). This is less obvious in the case of Alphazero, where the system is trained exclusively by playing against itself, based on a generic reinforcement learning algorithm. But, Alphazero has a very restricted and unambiguous task, even if the calculations it needs to perform are massive and complex; it does very well in a closed, rule-based game. In that sense it has no general artificial intelligence and cannot “survive” or even advise outside the confines of a rule-based game. See, e.g., Jose Camacho Colladoe, *Is AlphaZero Really a Scientific Breakthrough in AI?*, MEDIUM (Dec. 11, 2017), <https://medium.com/@josecamachocollados/is-alphazero-really-a-scientific-breakthrough-in-ai-bf66ae1c84f2> [<https://perma.cc/BKN4-NF2R>]; see also Rodney Brooks (@RodneyABrooks), TWITTER (Dec. 11, 2017 7:25 PM), <https://twitter.com/rodneymarbrooks/status/940784623574925312> [<https://perma.cc/6CMB-TYAP>].

the system can do without the acuity of human judgment, mistaking the imitation for what is imitated.⁴⁶ This is particularly concerning where platforms move into hardware, reconfiguring our built environment into a cyberphysical infrastructure that determines the “choice architecture” we face in everyday life—further erasing the difference between online and offline and further complicating our access to reality with a layer of backend systems that determine the hidden affordances of the interfaces. Google and Uber are moving into self-driving cars; Amazon has patented a wristband for its warehouse employees that enables monitoring their movements as well as prompting them into specific directions;⁴⁷ both are moving into smart homes and smart energy (Nesta, Alexa); Alphabet is moving into smart cities.⁴⁸ Thus, data-driven platforms are shaping the material world we navigate. In 1994, in the early days of the World Wide Web, Agre wrote a pioneering article about the way that computing systems ‘capture’ data,⁴⁹ implying a measure of violence as these systems force their environment (us) into the modus of data-engines, requiring a permanent translation of the flux of life into machine readable bits and bytes. In 2018, this translation is even more consequential. Due to their reliance on data-driven decision systems, platforms need ever more behavioral data and as they move into cyberphysical systems they will require a complete overhaul of existing infrastructure. Luciano Floridi notes in his book on the Fourth Revolution,⁵⁰ that in the practice of robotics the design of a robot’s environment (called the envelop) is as important as the design of the robot itself. Rodney Brooks observes that driverless cars will require a number of transformations of our cityscapes:

I think the popular interpretation is that driverless cars will simply replace cars with human drivers, one for one. I do

⁴⁶ See Federico Cabitza, *Breeding Electric Zebras in the Fields of Medicine*, ARXIV:1701.04077 (2017), <http://arxiv.org/abs/1701.04077> [<https://perma.cc/Q6DZ-27DX>]; Rich Caruana et al., *Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-day Readmission*, PROCEEDINGS OF THE 21ST ACM SIGKDD INTERNATIONAL CONFERENCE ON KNOWLEDGE DISCOVERY AND DATA MINING 1721 (2015) <http://doi.acm.org/10.1145/2783258.2788613> [<https://perma.cc/KE37-P2QE>].

⁴⁷ Olivia Solon, *Amazon Patents Wristband that Tracks Warehouse Workers’ Movements*, GUARDIAN (Jan. 31, 2018, 7:35 PM), <http://www.theguardian.com/technology/2018/jan/31/amazon-warehouse-wristband-tracking> [<https://perma.cc/BP48-MP6H>].

⁴⁸ Molly Sauter, *Google’s Guinea-Pig City*, THE ATLANTIC (Feb. 13, 2018), <https://www.theatlantic.com/technology/archive/2018/02/googles-guinea-pig-city/552932/> [<https://perma.cc/Q657-MG5S>].

⁴⁹ Agre, *supra* note 23.

⁵⁰ LUCIANO FLORIDI, *THE FOURTH REVOLUTION: HOW THE INFOSPHERE IS RESHAPING HUMAN REALITY* (2014)

not think that is going to happen at all. Instead our cities will be changed with special lanes for driverless cars, geofencing of where they can be and where cars driven by humans can be, a change in the norm for pick up and drop off location flexibility, changes to parking regulations, and in general all sorts of small incremental modifications to our cities.⁵¹

Whereas Brooks suggests that such transformations are incremental, I believe that together with similar developments in energy grids, drone delivery, retail, healthcare, education and employment, our cityscapes, homes, offices, hospitals, schools, shops and factories will undergo major reinvention and reconfiguration. Cyberphysical infrastructures will require a major overhaul of our shared environment to ensure physical safety and to secure their functionality. This will entail both material scaffolding to protect human beings and further datafication to “optimize” the utility of these systems (at least for those who invest in them). As these systems will be driven by unprecedented amounts of machine-to-machine communications, while developing new types of dynamics once they start interacting (smart energy grids “talking” to connected electric cars; digital assistants for smart homes “exchanging notes” with each other in the cloud; artificial legal intelligence systems of law firms interacting with those of courts, public prosecutors and other law firms), we may expect a reduction of the interstices where semantic discontinuity would have developed in a text-driven environment. This does not necessarily mean that semantic discontinuity will be eradicated altogether; as long as we continue to speak, write, and read there will be ambiguity and a poly-semantics that is open to competing interpretation. However, as our material and institutional environment is captured, modulated, and reconfigured by interacting data-driven systems, the impact of our musings may be inconsequential. We will have to anticipate how these systems “read” us and adapt our behavior to achieve the goals that their architecture allows us to attain, as they will determine our choice architecture. This will fundamentally transform our understanding of both ourselves and our environment, as it will require us to partake in the data-driven closure that is inherent in the computationally informed environment that we need to navigate. The option to resist, to dissent, and to develop dissonance will depend on a contrapuntal movement that must fit the logic of the cognitive machines that will populate our environment.

⁵¹ Rodney Brooks, *My Dated Predictions*, RODNEY BROOKS: BLOG (Jan. 1, 2018), <https://rodneybrooks.com/my-dated-predictions/> [https://perma.cc/3H6K-HTGR].

It is, however, not obvious that such cognitive machines generate the semantic discontinuity that affords dissonance in the first place.

III. THE PRIMITIVES OF LEGAL PROTECTION IN A DATA-DRIVEN PLATFORM ECONOMY

The issues discussed above cannot be solved in one stroke, as cyberphysical architectures are already moving into the capillaries of our societal architecture. A potent cocktail of effective competition law, consumer law, tort law, global tax law, human rights law, and data protection law will have to be negotiated to reduce the capture of data-driven platforms over our onlife world. We need to ensure, however, that the smallest building blocks or primitives of legal protection must be situated at the level of the data ecosystem itself, capable of addressing automated, data-driven decisions (whether based on machine learning or on predefined parameters, as in smart contracts). Speaking of primitives highlights their pre-conditional character, as they should offer protection at the gateways of the data-driven ecosystem, preventing a host of problems that must otherwise be solved at other levels, where the cost of both compliance and enforcement would be much higher. In the context of digital security such a primitive can be found in the adage “select before you collect.” In the generic context of data ecosystems, three interlocking primitives of legal protection can be detected: data minimization, purpose limitation, and justification of automated decisions.

A. Data Minimization and Purpose Limitation

Data minimization means that only those data are processed (collected, stored, mined, inferred, used for training algorithms) that are necessary. Processing of any excess data is unnecessary, thereby creating unnecessary risks. These risks vary from hacking to unreliable inferences resulting in incorrect, wrongful, and potentially dangerous decisions. As decisions are often based on complex machine-to-machine exchanges, high-dimensional neural networks, and unpredictable dynamics between interacting systems, it is better to prevent training on data that is not necessary. This then raises the questions of what counts as necessary, who gets to decide this, and in the light of what purpose the data is necessary. What person, technology developer, provider, regulator or legislator has the power of definition here?

Let's first establish that the concept of necessity is not a “standalone device.” Data can be necessary for one purpose but not for another, which means the data can be processed for that one purpose but not for the other. In the domain of data protection, many lay persons have

come to believe that any data can be of interest at some time in the future and may thus become necessary. This is a truism. It does not, however, imply that every data is necessary anytime for whatever purpose. On the contrary, from the perspective of the science of machine learning, data that is irrelevant, incorrect, or incomplete is not merely redundant but will give rise to bad output. For this reason Van der Lei formulated his first law of informatics: “Data shall be used only for the purpose for which they were collected. This law has a collateral: If no purpose was defined prior to the collection of the data, then the data should not be used.”⁵² This confirms that data minimization is not just a matter of data protection law, but a precondition for the methodological integrity and robust, reliable output of machine learning. Data minimization will reduce the need to curate or cleanse data, but it will also prevent working with low-hanging fruit, which is usually not relevant, let alone necessary.

The General Data Protection Regulation (GDPR), applicable as of May 25, 2018,⁵³ requires that—to be lawful—processing of personal data must be necessary for at least one of five legal grounds: contract, vital interests of the data subject, a legal obligation, a public task, or the legitimate interest of the data controller.⁵⁴ Alternatively, the processing can be based on unambiguous and informed consent.⁵⁵ Such consent, however, depends on one or more explicitly specified purposes and the processing must—again—be necessary for such purpose(s).⁵⁶ On top of that, consent can be withdrawn at any time and the withdrawal must be as easy as the provision of consent.⁵⁷ Once consent is exhausted (e.g., because the purpose has been achieved) or withdrawn, the processing must stop (for which an injunction can be filed in a court of law).⁵⁸ Additionally, recital 43 explains that “[c]onsent is presumed not to be freely given if . . . the performance of a contract, including the provision of a service, is dependent on the consent despite such consent not being necessary for such performance.”⁵⁹ This basically prohibits service

⁵² J. van der Lei, *Use and Abuse of Computer-Stored Medical Records*, 30 METHODS OF INFO. IN MED. 79–80 (1991); Federico Cabitza et al., *A Giant with Feet of Clay: On the Validity of the Data that Feed Machine Learning in Medicine*, ARXIV:1706.06838 [CS, STAT] (2017), <http://arxiv.org/abs/1706.06838> [<https://perma.cc/CVT7-2FZB>].

⁵³ Council Regulation 2016/679 of the European Parliament and of the Council on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Advancement of Such Data, and Repealing Directive 95/46/EC, 2016 O.J. (L 119) 1 [hereinafter the General Data Protection Regulations or GDPR].

⁵⁴ *Id.* at art. 6(1)(b)–(f).

⁵⁵ *Id.* at art. 6(1)(a).

⁵⁶ *Id.* at art. 5(1)(c).

⁵⁷ *Id.* at art. 7(3)–(4).

⁵⁸ *Id.* at art. 79.

⁵⁹ *Id.* at art. 7(4).

providers from demanding consent for the processing of personal data that is not necessary for a service, if the provision of that service is made conditional on consent for such additional processing (which, e.g., implies a tracking-wall prohibition).

The right to data portability enables a request “to have the personal data transmitted directly from one controller to another, where technically feasible⁶⁰ . . . in a structured, commonly used and machine-readable format.”⁶¹ If properly enforced (by supervisors and courts), the combination of the right to withdraw consent at any time, as easy as it was provided, with the right to data portability, will reconfigure the market for personal data ecosystems, uproot business models based on excessive pricing, and reduce platforms’ ability to do whatever they see fit in terms of data capture.

All eyes are now on enforcement: by the national data protection authorities,⁶² by national courts,⁶³ and by the Court of Justice of the European Union. Interestingly, the GDPR requires that platforms get their act together in terms of compliance at the backend of their systems⁶⁴ by stipulating demonstrable data protection by design⁶⁵ and proper documentation of all relevant forms of processing,⁶⁶ while providing data protection authorities with far reaching powers to investigate compliance.⁶⁷

Instead of (1) assuming that the American way of life will prevail as far as data protection is concerned, and (2) assuming that data minimization and purpose limitation are bad for innovation in machine learning, I conclude that the EU approach has the potential to redefine the market for data ecosystems in a way that is (a) conducive to the semantic discontinuity that enables generative semantic dissonance and may (b) save us all from bad innovation and from dangerous dependence on data-

⁶⁰ Article 20 speaks of data “provided by” the data subject. The Article 29 Working Party has interpreted this as referring to both volunteered and observed data, even if inferred data are not covered. See *Article 29 Data Protection Working Party, Guidelines on the Right to Data Portability*, at 9–10, WP 242, 16/EN (Dec. 13, 2016). See also Paul de Hert et al., *The Right to Data Portability in the GDPR: Towards User-Centric Interoperability of Digital Services*, 34 *COMPUTER L. & SECURITY REV.* 193 (2018).

⁶¹ GDPR, *supra* note 53, at art. 20(1).

⁶² *Id.* at art. 77–78, 80, 83.

⁶³ *Id.* at art. 78–80, 82.

⁶⁴ Sheera Frenkel, *Tech Giants Brace for Europe’s New Data Privacy Rules*, *N.Y. TIMES* (Jan. 28, 2018), <https://www.nytimes.com/2018/01/28/technology/europe-data-privacy-rules.html> [<https://perma.cc/J42U-EQKG>].

⁶⁵ GDPR, *supra* note 53, at art. 25.

⁶⁶ *Id.* at art. 30.

⁶⁷ See *id.* at art. 58.

driven critical infrastructures that are flawed due to incorrect, irrelevant, or incomplete data.

B. Justification of Automated Decisions

However, data minimization and purpose limitation concern legal protection regarding the input of machine learning (ML) systems. As primitives of legal protection, they do not protect against bad machine learning research designs, nor against unjustifiable automated decisions. Here, what is needed is (1) a set of standards that testifies to the methodological integrity of exploratory or confirmatory ML research designs,⁶⁸ and (2), a clear focus on the justification of automated or semi-automated decisions.

Good ML depends on a set of design decisions, each of which has trade-offs. First, we have decisions about which data to collect and how (via sensors, tracking of click-stream behaviors, surveys), how to curate the data (“cleanse,” as they say) and how to divide it into training and validation sets. Trade-offs include the costs and the speed of acquiring the data (does one resist the temptation to go for “low hanging fruit?”), and the completeness of the data (is it viable to construct synthetic data if the data is incomplete or will blind spots appear anyway, making the whole exercise hazardous, depending on the purpose?). The second stage concerns the design of the hypotheses space that determines how algorithms will be trained on the data. Note that the hypotheses are mathematical functions which developers hope will approximate an assumed target function. This again involves a series of trade-offs, as different types of algorithms may affect predictive accuracy, the explainability, and the out-of-sample reliability in different ways. The second stage—which should overlap with the first as both types of design decisions interact—also involves the choice of the task for which the algorithms are trained (even in the case of unsupervised ML a task must be formulated, otherwise the system cannot learn) and the performance metric. Tasks and performance metrics make all the difference for the “predictive accuracy” that is often heralded as a key indicator for systems that are said to “outperform” human experts. Depending on which performance metrics are used, the predictive accuracy will differ, though this may not at all be obvious to those buying, using, or being targeted by a system. As predictive accuracy can only refer to the way the algorithms predicts historical output data (nobody can train a system on future data), confirmatory research design requires extensive testing on out-of-sample

⁶⁸ Jake M. Hofman et al., *Prediction and Explanation in Social Systems*, 355 SCIENCE 486, 486–88 (2017).

data. And even then, high accuracy in the data may not make any sense for the reality outside the data set, leading to major consequences if this is not acknowledged or if this remains invisible.⁶⁹

If automated or semi-automated decisions (where human involvement is merely “fabricated”)⁷⁰ seriously affect a person, she should be made aware of the automated nature of the decision and provided with sufficiently understandable information to be able to either change her behavior or contest the decision.⁷¹ Basically, she should be able to halt the making of such decisions, if they cannot be justified.⁷² This is crucial if we want to safeguard the interstices of semantic discontinuity that enable dissent and dissonance. Though it is important that decisions of automated systems can be explained (whether ex ante or ex post; whether individually or at a generic level), we must keep in mind that in the end what counts is whether such decisions can be justified.⁷³ An explanation of the system’s logic and its envisaged consequences is not necessarily enough to justify its decision.

Depending on whether an automated decision has been taken by a government agency, a not-for-profit organization, or a private company, different legal regimes will apply as to justification. Justification of a decision cannot be based on the GDPR, which only concerns the justification of a specific aspect of the decision. For public administration decisions the legality principle requires legal competence and, depending

⁶⁹ Caruana et al., *supra* note 46; Federico Cabitza, *The Unintended Consequences of Chasing Electric Zebras*, https://www.researchgate.net/publication/311702431_The_Unintended_Consequences_of_Chasing_Electric_Zebra [https://perma.cc/8Q4F-VPDK].

⁷⁰ Article 29 Data Protection Working Party, *Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679*, at 10, WP 251 16/EN (Oct. 3, 2017).

⁷¹ GDPR, *supra* note 53, at art. 13(2)(f), 14(1)(g), 15(1)(h) (providing an obligation to provide and a right to obtain such information).

⁷² *Id.* at art. 22 (providing a qualified right not to be subject to such decisions, with relevant exceptions, enforceable under art. 79).

⁷³ An entire literature has sprung up on this point. *See, e.g.*, Bryce Goodman & Seth Flaxman, *European Union Regulations on Algorithmic Decision-Making and a “Right to Explanation,”* ARXIV:1606.08813 [CS, STAT] (2016), <http://arxiv.org/abs/1606.08813> [https://perma.cc/3AES-JYXH]; Sandra Wachter et al., *Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation*, 7 INT’L DATA PRIVACY L. 76 (2017); Lilian Edwards & Michael Veale, *Slave to the Algorithm? Why a “Right to Explanation” is Probably Not the Remedy You are Looking For*, 16 DUKE L. & TECH. REV. 18 (2017); Gianclaudio Malgieri & Giovanni Comandé, *Why a Right to Legibility of Automated Decision-Making Exists in the General Data Protection Regulation*, 7 INT’L DATA PRIVACY L. 243 (2017); Andrew Selbst & Julia Powles, *Meaningful Information and the Right to Explanation*, 7 INT’L DATA PRIVACY L. 233 (2017).

on the applicable legal norms, a test of their lawfulness—while considering that decisions may also infringe human rights, thereby requiring additional testing (e.g., the so-called triple test of art. 8-11 European Convention of Human Rights, concerning (1) a legitimate aim, (2) lawfulness, (3) necessity (in the sense of proportionality). Decisions made by private corporations may not require a justification because the company may claim it is free to contract as it sees fit. In the context of the EU, this appeal to the freedom of contract can be reframed as an appeal to the fundamental freedom to conduct a business.⁷⁴ Clearly this fundamental right is not absolute, meaning that relevant restrictions apply, such as those against unfair commercial practices;⁷⁵ unfair contract terms;⁷⁶ discrimination based on race and ethnic origin;⁷⁷ employment discrimination on grounds of religion or belief, disability, age, or sexual orientation;⁷⁸ restrictions that aim to prevent violations of equal treatment for men and women in matters of employment and occupation;⁷⁹ and equal treatment for men and women in the access to and supply of goods and services.⁸⁰ Note that the concept of indirect discrimination (particularly relevant in the case of machine learning) may trigger the applicability of non-discrimination law, thus placing a major restriction upon the freedom to conduct a business (which encompasses, e.g., the freedom to contract and the freedom to develop a competitive business model). Clearly, the GDPR obligation to provide meaningful information about the logic of processing should enable individuals or their representatives⁸¹ to obtain enough information to detect indirect discrimination.⁸² Other limitations of the freedom to conduct a business will apply (based on other human rights, environmental law, etc.). The point is to make sure that the discussion about explainable AI does not displace the discussion about the need to justify an automated decision. The mere fact that a decision is deemed expedient or accurate (in terms of what?) by either a government

⁷⁴ Charter of Fundamental Rights of the European Union art. 16, Dec. 18, 2000, 2000 O.J. (C 364) 12.

⁷⁵ Council Directive 2005/29/EC, 2005 O.J. (L 149) 22.

⁷⁶ Council Directive 93/13/EEC, 1993 O.J. (L 95) 29.

⁷⁷ Council Directive 2000/43/EC, 2000 O.J. (L 180) 22.

⁷⁸ Council Directive 2000/78/EC, 2000 O.J. (L 303) 16.

⁷⁹ Council Directive 2006/54/EC, 2006 O.J. (L 204) 23.

⁸⁰ Council Directive 2004/113/EC, 2004 O.J. (L 373) 37.

⁸¹ GDPR, *supra* note 53, at art. 79, 80.

⁸² On the concept of indirect discrimination in the case of the both European Courts, see Emilie Delcher, *Overview of the Case Law on the Prohibition of Discrimination of the ECJ and ECtHR*, ICELANDIC HUMAN RIGHTS CTR. (2012), <http://www.humanrights.is/en/activities/equality-non-discrimination/overview-of-the-case-law-on-the-prohibition-of-discrimination-of-the-ecj-and-ecthr> [https://perma.cc/4TKM-2WP8].

agency or a private corporation is not a sufficient reason to justify a decision. It will be more interesting to develop actionable information about automated decisions for potential justification.⁸³ The legal primitives here are (1) the default right not to be subject to automated decisions that seriously affect data subjects and (2) the legal obligation to provide relevant information, as this will often be a precondition to contest its justifiability.

The combination of an enforceable right not to be subject to unlawful automation and the right to obtain relevant information, together with the right to mandate such rights via collective action, should reconfigure the dependence of markets, business models, and administrative decision-making on automated systems—including the reconfiguration of the backend systems of data-driven platforms. To preserve and reinvent spaces of semantic discontinuity, competition law, consumer law, and tort law will not suffice. Their impact will co-depend on safeguarding a set of primitives of legal protection that are tuned to a major reform of current data ecosystems.

⁸³ See, e.g., Sandra Wachter et al., *Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR*, 31 HARV. J. L & TECH (forthcoming 2018) <https://papers.ssrn.com/abstract=3063289> [<https://perma.cc/Z25X-5PKR>].