

# Artificial Intelligence and the Purpose of Social Systems

Sebastian Benthall  
New York University School of Law  
USA  
spb413@nyu.edu

Jake Goldenfein  
Melbourne Law School  
Australia  
jakeg@unimelb.edu.au

## ABSTRACT

The law and ethics of Western democratic states have their basis in liberalism. This extends to regulation and ethical discussion of technology and businesses doing data processing. Liberalism relies on the privacy and autonomy of individuals, their ordering through a public market, and, more recently, a measure of equality guaranteed by the state. We argue that these forms of regulation and ethical analysis are largely incompatible with the techno-political and techno-economic dimensions of artificial intelligence. By analyzing liberal regulatory solutions in the form of privacy and data protection, regulation of public markets, and fairness in AI, we expose how the data economy and artificial intelligence have transcended liberal legal imagination. Organizations use artificial intelligence to exceed the bounded rationality of individuals and each other. This has led to the private consolidation of markets and an unequal hierarchy of control operating mainly for the purpose of shareholder value. An artificial intelligence will be only as ethical as the purpose of the social system that operates it. Inspired by the science of artificial life as an alternative to artificial intelligence, we consider data intermediaries: sociotechnical systems composed of individuals associated around collectively pursued purposes. An attention cooperative, that prioritizes its incoming and outgoing data flows, is one model of a social system that could form and maintain its own autonomous purpose.

## CCS CONCEPTS

• **Social and professional topics** → **Computing / technology policy; Computing / technology policy**; • **Security and privacy** → **Human and societal aspects of security and privacy**; • **Computing methodologies** → **Philosophical/theoretical foundations of artificial intelligence**.

## KEYWORDS

artificial intelligence, cybernetics, economics, intermediaries, liberalism, platforms, privacy

### ACM Reference Format:

Sebastian Benthall and Jake Goldenfein. 2021. Artificial Intelligence and the Purpose of Social Systems. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*, May 19–21, 2021, Virtual Event, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3461702.3462526>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

AIES '21, May 19–21, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8473-5/21/05...\$15.00

<https://doi.org/10.1145/3461702.3462526>

The law and ethics of Western democratic states have their basis in liberalism, and this extends to legal regulation and ethical discussion of technology and businesses doing data processing. Liberalism relies on the privacy and autonomy of individuals, their ordering through a public market, and, more recently, a measure of equality guaranteed by the state. We argue that these forms of regulation and ethical analysis are largely incompatible with the techno-political and techno-economic dimensions of artificial intelligence. Over the course of the twentieth century, computer science, cognitive psychology, operations research, management science, statistics and other fields, have converged on a science of agency. This science explains and enables the supremacy of corporate actors with artificial intelligence capabilities over individual humans that often lack the capacity to know and defend their own interests. Regulatory solutions that seek to reinforce liberal patterns of individual privacy and public markets have enabled the data economy, which is different from anything imagined in liberal legal theory, to develop further along this trajectory. For instance, platforms have inverted the relationship between individuals and the market, making the former public and the latter private.

We see this disconnect between liberalism as a basis for ethical reasoning in Western legal systems and artificial intelligence in fact as the source of anxiety around the power of large technology companies. Artificial intelligence's challenge to ethics is its challenge to liberal political theory. In Section 1, we trace liberalism's approaches to rendering technology ethical through privacy laws, regulation of the public market, and pursuit of "fairness in AI". In our analysis, these liberal interventions have not entirely succeeded. Rather, the data economy and artificial intelligence have transcended legal imagination. More recent scientific advances, especially those connected to the theory of the firm, have revealed liberalism's blind spots.

In Section 2, we consider artificial intelligence for itself. The term "artificial intelligence" is imprecise. Originating in a highly influential workshop at Dartmouth College in 1956, it can now mean variously a research field within computer science, a class of technologies, or kind of product or service. Generally the term is used to designate the cutting edge of what is possible using computing technology; this is captured in the quip or joke known as Tesler's Theorem: "AI is whatever hasn't been done yet." Originally, the research questions around artificial intelligence centered on the question of whether a machine could think like a human being [104]. This question of psychological verisimilitude has since given way to the question of how to use computing technology to perform complex information processing tasks, such as image classification, language translation, and transcription of text into speech, at increasing scale, through statistical and optimization techniques known as machine learning. Academic research communities still develop "artificial intelligence" in ways that are decontextualized

from industrial processes. However, as the term has come to be used to refer more and more to commercial products and services, it must be understood to denote a much more complex system of user interfaces, business to business exchanges of data and computing services, and contracted, sometimes invisible, labor [48].

Early founders of the field of artificial intelligence such as Herbert Simon [87] understood that the “sciences of the artificial” were a combined science of computing machines, information systems broadly construed, and organizational management. This scientific paradigm has led to profound social shifts. These include the reconceptualization of the individual human in terms of information processing operations that can be compared to those of a machine. In the commercial context, many of the operations typically attributed to the entrepreneur (such as search and forecasting) have now been automated, making the firm more subject to absentee ownership. Artificial intelligence has enabled frictionless, “smart”, data-driven, and privately owned alternatives to the public price mechanism via the platform as an economic formation. When individual humans can be compared directly, as agents, with larger socio-technical organizations such as firms, the fact that rationality is bounded, and differently bounded, is a fundamental inequality that flouts any egalitarian principle. These scientific facts and economic realities have upended liberal assumptions. An applicable, actionable normative theory must build on the theory of artificial intelligence, not be blind to it.

In Section 3, we explore one way to emerge from this crisis. Atomized individuals will not be able to guarantee their interests and freedoms within the privately controlled superordinate agencies of platforms that treat them as “users”. Instead, ethical AI must enable collective agency through new data and sociotechnical infrastructure designed for different purposes. We look to the scientific paradigms of systems theory, such as the social theory of Niklas Luhmann, and the technical discipline of artificial life as foundations for this new line of inquiry. If liberalism is limited to the understanding of society divided into public and private spheres, artificial intelligence ethics requires an understanding of society composed of multiple, intersecting social systems with diverse purposes [71].

We review several past and proposed data intermediaries that are examples of social systems in this sense. These systems are designed for various purposes: greater individual control of data, outsourced policing and enforcement of individual rights, better governance of information flows, and the remission of value to consumers, are just a few. The vast majority of data intermediaries fail. We distinguish between those intermediaries designed according to liberal ideas, and those that instead are aimed at collective purposes.

We come to the conclusion that *an artificial intelligence will be as ethical as the purpose of the social system that operates it*. Hence, an AI that is operated by an absentee-owned corporation that treats individual users as means, not ends, will never be ethical. Rather, the law and ethics of AI must imagine new forms of sociotechnical collectives that use “smart” technologies to further the purposes of individuals as members. Ethical AI is artificial intelligence used for self-governance, the coordination of authentic collective attention and action, and sustainable life.

## 1 LIBERALISM AND THE DIGITAL

This section describes certain juridical features of liberal political theory that condition existing strategies for regulating artificial intelligence practices and business models. These include, first, an individual person, fashioned as morally autonomous [103] and self-determining [96], that is endowed with sufficient rationality to exercise their rights freely and effectively. Second, the privileging of individual autonomy conditioned by both privacy (as a zone of freedom from state coercion and other private parties) [38] and the capacity to possess and alienate private property [72]. In certain Lockean liberal traditions, because the individual is self-possessing, the capacity to alienate private property extends to the capacity to alienate oneself both in labor and in image [34]. Third, this equal [112], universal, and autonomous subject is capable of transacting and exchanging its rights with other equals through a free public market [53]. Finally, and most recently, some branches of liberal theory have taken on egalitarianism as a core principle [5, 76].

Liberalism imagines a society of individuals. These individuals are endowed with moral rationality and autonomy. They coordinate by means of the market, a public system of exchanges of what is otherwise private—private property. Liberally premised law and ethics governing artificial intelligence therefore aim for privacy and data protection to preserve individual autonomy, better markets to preserve data as private property, and interventions on machine learnt models to ensure egalitarian fairness. We examine how these liberal interventions into artificial intelligence applications have fared in practice and find that they are failing because they unfold on a political and theoretical terrain that has been superseded by artificial intelligence itself.

### 1.1 The individual subject and its autonomy through privacy and property

Liberal law presumes an individual that is an independent center of consciousness with agency and the capacity to be held responsible for action. External direction or imposition as to what the true interests of the individual might, or ought to, be is thus an improper interference with both dignity and autonomy. The individual is the cornerstone of liberal moral reasoning. The freedom to choose is paramount, and the capacity to choose is expressed, on one hand, through the possession of rights and freedom to contract, and on the other, by the capacity to author one’s own identity. Private property is the name given to the material part of the protected individual domain [53], the power of which is expressed in the legal capacity to enforce those property rights against the world at large. In liberalism, “privacy” protects the immaterial dimensions of the individual domain against both the state—as a solution to preserving individual freedom while simultaneously legitimating sovereign power (i.e. the rule of law)—and against private actors.

Where to draw the line between the legitimate and illegitimate domain of the state (i.e. the extent of privacy and the comprehensiveness or limits of private property rights), and on what ethical or moral premise, are both debated. However, all shades of liberal thought rationalize the non-state-governed dimensions of life into rights, especially property rights, ordered by private law. Through property rights, it becomes possible to govern groups of individuals that do not share similar goals, purposes, or notions of the good.

In the digital economy, the individual domain is primarily protected through privacy and data protection laws. Both “dignitarian” (including ideas around subject self-creation, see [30, 31]) and “control” (or “liberty” [114]) based justifications for privacy can be understood as defending a sphere of individual autonomy from illegitimate interference. Governing the movement and processing of data through regulations premised on individual economic “control” [113] typically relies on “notice and consent” to define the parameters of legitimate action. That is, autonomy is protected because individuals as “data subjects” consent to the collection and processing of their personal information. The notice and consent paradigm is widely acknowledged to be a failure [9]. However it remains dominant in existing and emerging data protection legislation. Dignitarian approaches in privacy typically define elements of selfhood as non-trafficable, making the exercise of data protection and privacy rights an act of self-authorship. Deriving from Kantian ideals in European constitutional protections of “personality” [63], ideas like “informational self-determination”<sup>1</sup> and “information sovereignty” have become central to European data protection law. Control over data is thus a means of self-presentation and self-authorship. Rather than contractual economic exchange of data for a digital service, dignitarian rationales enable the individual to define how their identity is represented in databases controlled by external entities like states and companies. Although European data protection law does *not* rely on consent as the exclusive mechanism for legitimate data processing, consent and individual control over data (in the form of “data subject rights” that enable access, rectification and erasure of data, as well as the contestation of profiles) remain the primary instruments for regulating online behavioural advertising. Indeed, one focus of the European Data Protection Regulation (GDPR), likely the strictest data protection regulation globally, was to upgrade consent requirements to “meaningful consent”. This manifests as endless “popups” performing transparency about the movement and processing of personal data. In the world of behavioural advertising and the digital economy more broadly, the dignitarian protections of the European General Data Protection Regulation (GDPR) are thus reconcilable with the more individual economic “control” oriented governance of the movement and processing of data in more property-oriented liberal jurisdictions. Whatever its philosophical or political foundation, this form of individual control over personal data, either as an economic resource or elements of personality [110], presents little barrier to the extraction and monetization of personal data through the generation and trading of inferences and insights, which typically elude data protection regulation [55].

In the AI informed digital economy, however, data is better understood as a means of governance—not simply a commodity or resource [19, 73]. In other words, data is meaningful because, like an electrical charge, it flows in a circuit from and then back to the individual, establishing a dynamic, ongoing relationship. Between its beginning and end, the data flows between unaccountable intermediaries, each extracting value from the control system. Data becomes a tool for governing others through its capacity to be refined into predictive products that, in this digital economy, are

used to steer consumer behaviour. The liberal representational account of data as static and transactional attaches awkwardly to the reality of data, co-generated by individuals through platform interactions, and channeled to and from networked actors on the other side of the platform who use it in real-time. Data is also no longer merely personal because it inevitably reveals information about multiple individuals [10, 37, 73, 110], and has value mainly through its aggregation into audiences and consumer segments. Because social information derived from data does not pertain directly to an individual, it becomes an externality in the liberal market conception [54], and beyond the scope of liberal regulation. As we discuss below, a more engaged form of regulation would move on from individual control over the inputs of those aggregated data products to the establishment of autonomy, agency, or control over the outputs of data processing [37]. Individual control over data does not translate into individual control over the governing power of data. In other words, self-authorship does not translate to self-governance in the digital world.

## 1.2 A free public market?

Although some enthusiastically distinguish between liberalism and neoliberalism, or classical liberalism and social liberalism, or liberalism and libertarianism [41], the different governing forms embodying these political theories only shift the degree of freedom or autonomy guaranteed from the state (or alternatively the scope of limited government), and what elements of life are placed in market domains [53]. While their social consequences clearly vary, they share certain key organizational ideas. Indeed, for Hayek, the essence of liberalism is a self-generating or spontaneous ordering in social affairs. The spontaneous order takes the form of unregulated or free markets that are the inverse of a coordinated or planned ordering that is the realm of democratic governance. Noting that few markets are actually free or unregulated, the complex forms of governance that have developed through the 20th and early 21st century have still enhanced the role of markets in coordinating social life through legal tools that instantiate and then protect those markets from democratic intervention [23, 89].

Liberal scholars, originating with Laudon [65], have proposed that it would be beneficial to organize the digital distribution of personal data directly using public markets. This proposal has faced many practical and theoretical challenges. A primary difficulty of instituting personal data markets is the creation of a viable property right to personal data. Data is naturally non-rivalrous and easy to copy; a data market would need to establish its alienability, rivalry, and excludability somehow [90]. While it is possible to imagine something akin to an intellectual property right to personal data, such a regime would be inadequate because of the intrinsic relationship between personal data and privacy, which is also valued by liberalism. Some argue that whereas intellectual property is alienable, privacy is not (or ought not be); ergo, intellectual property rights in personal data would lead to a contradiction within liberal legal theory [79], at least in its dignitarian formulations. Others have advocated for treating personal data creation as a form of labor, which would then be sold on a market [7, 64]. Perhaps labor is justly alienable from the laborer, and perhaps not—either way,

<sup>1</sup> German Census Act Case, Bundesverfassungsgericht [German Constitutional Court], 1 BvL 209, 269, 362, 420, 440 and 484/83, 15 December 1983 reported in (1983) 65 BVerfGE 1.

this recommendation still runs into the well-known objections to data markets.

A further critique of data markets and the corollary of formalized property rights in data (as well as the idea of data as labor), is that this seems unlikely to alter the political economy of data processing [109]. Property rights will not disturb the dynamics of a regulatory system that already treats data as a pseudo-commodity, protected by trade secrets, technical infrastructures, and contracts that enable the enclosure and primitive accumulation of data by platforms [32]. Platforms are more like firms than markets, while at the same time superseding the liberal imagination of the firm.

Economics perennially debates the nature of the firm. Marshall, in his *Principles of Economics* (1890), noted organization as a fourth “factor of production”, along with land, labor, and capital. Coase (1937/1995) would later famously argue that the purpose of the firm is to allow the entrepreneur to purchase factors of production without the “costs of using the price mechanism”, which include the costs of discovering prices, negotiating, and managing risks. Coase’s transaction cost economics also suggests the existence of a type of organizational rationality governing firm decision making, rather than the firm existing as merely an aggregation of individual agents seeking to maximize their own self interest as they conduct organizational business [119]. That is, the firm has its own organizational utility of profitability, which may not always translate into the interests of the firm’s members. And while certain economists have proposed to define corporate firms, for instance, as a nexus of contracts [22], the firm maintains a hierarchy which distributes information and power unequally amongst those individuals. Williamson [115] has elaborated transaction cost economics to account for the specific conditions—such as asset specificity, environmental uncertainty, and the threat of opportunism—under which it is inefficient for the firm to use the market rather than develop resources under its hierarchical control. These conditions are pervasive in the digital economy, in which personal data and artificial intelligence technologies are specific assets for managing uncertainty. Now firms commonly offer *ex ante* “free” services that bind users into contracts with unknown *ex post* costs [58]. The employees of these firms are similarly bound by nondisclosure agreements into privately ordered arrangements in which their ability to act as public citizens is limited. Corporate organizations should not be thought of as collectives in a meaningful sense—they are their own artificial entities.

In this framing, the firm and the entrepreneur are interchangeable, while employees are simply factors of production coordinated and bound by contracts into a production function. The entrepreneur constructs the firm out of contracts to reduce the costs of search, negotiation, and forecasting as well as their exposure to risk. However, in the mid-20th century scientists began mastering and automating these tasks of managing information costs. The entrepreneur, or fourth factor of production of “organization”, became a kind of technology: artificial intelligence. The firm as platform uses its artificial intelligence to automate and coordinate, in service of its entrepreneurial profit-goal, its internal activities as well as those of its external users. Through automation and experimentation at scale, the firm as platform manipulates the visibility, rankings, and prices of their users’ relations and transactions with each other. As these deviations from ideal market conditions

become more apparent and egregious, regulators have sharpened another regulatory tool intended to preserve the competitiveness of markets—competition and antitrust law.

Liberalism requires of its entrepreneurs a modicum of fair play to maintain the benefits of competition. Regulators worldwide have responded with redoubled interest in antitrust law enforcement (i.e. with US regulator activity against Google and Facebook), legislative reforms targeting, *ex-ante*, platforms’ abuse of market dominance (i.e. with the EU Digital Services Act package), and new rules addressing transparency and financial flows between platforms and content industries (i.e. with the Australian competition regulator’s Treasury Laws Amendment (News Media and Digital Platforms Mandatory Bargaining Code) Bill 2020 (Cth)). Disciplining the anti-competitive practices of firms represents another liberal regulatory approach premised on the belief that healthy market competition generates better outcomes for consumers. Research has clearly demonstrated the various ways platforms abuse their dominant market positions across all sectors of the digital economy [91, 92]. Without closely scrutinizing these regulatory efforts here, we agree with Pistor [73] that it matters little to a data subject whether a monopoly incumbent or a startup is collecting and processing data because “a consumer who is the target of data’s predictive power does not necessarily benefit from greater competition among data controllers”. Ultimately, such interventions designed around the logic of liberal public markets have limited efficacy because they have not taken into account how platform environments have superseded both firms and markets as the primary form of economic organization in the digital economy [32]. Rahman [75] and others have proposed more radical interventions in the form of public utility-type regulations for platforms to improve transparency, neutrality, and rate-setting. But short of options achieving pseudo-socialization, these are unlikely to constrain business models or reduce abusive data practices because they continue to rely on markets mechanisms to coordinate service provision through a profit incentive. Without updated understandings of the way data, firms, users, and “markets” function in the digital economy, these approaches fail to tackle incentives to treat user data as a raw economic resource, or limit the exploitation of information flows through the infrastructures that platforms control [15].

### 1.3 Egalitarianism and fairness in AI

These previous motifs—the autonomous individual and the public market—are typically attributed to Classical Liberalism, and have each been subjected to critique from both inside and outside liberal thought for more than a century. Reforms have recognized the limits of individual autonomy and alienability, and power differentials in the market (i.e. the coercive capacities of individuals [51]). In contrast to Classical Liberalism, New Liberalism is characterized by liberal “justice” interventions [5]. The conditions and aims of equality are a topic of debate among contemporary liberal scholars, but perhaps are most paradigmatically expressed by Rawls in *A Theory of Justice* and later work. These include, for instance, human rights projects like civil and political rights, workers’ rights, and so on, as well as property laws that recognize the social dimensions of private property on diginitarian, civic, and social justice grounds. Unfettered Classical Liberalism is certainly no longer the norm

and more “social” forms of liberalism that recognize and address inequalities of distribution of social goods as well as relational inequalities have diminished the private autonomous sphere and increased the scope of legitimate regulatory intervention.

The egalitarian impulse of New Liberalism is very much alive in AI ethics under the key and contested word “fairness”. Without regulation, artificial intelligence applications risk exacerbating inequality through “rational discrimination” of people based on their unequal life circumstances [44]. While the Rawlsian principle of equal opportunity is most directly invoked in work on fairness in machine learning systems involved in the hiring of employees and compliance with anti-discrimination law [52, 60], the field has been energized by use cases in the field of criminal justice [16, 26]. These interventions into machine learning research were accompanied by the recognition that there are many different meanings of “fairness”, implying distinct tradeoffs and requiring sensitivity to context and worldview [42, 62, 118]. Academic debate about fairness in AI has sometimes been heated, as critical scholars have advocated that fair or egalitarian ends require an understanding of the economic and social contexts in which AI operates [49, 82], including its relationship with technology company hiring practices [88]. Some scholars have argued that in many cases ethics implies the refusal of the use of AI at all [28, 74, 93]. Despite all the efforts to align artificial intelligence applications with the egalitarian ethos, the fundamental inequality that remains is between societies of individuals and the powerful firms that own and operate artificial intelligence. Egalitarianism focused on distribution between individuals has thus elided thorough examination of the nature of artificial intelligence as a manifestation of economic organization.

## 2 ARTIFICIAL INTELLIGENCE FOR ITSELF

In liberal jurisdictions, the social frictions surrounding artificial intelligence have been understood through liberalism’s conceptual frames: the individual rational subject, their autonomy and privacy, their coordination through the public market, and the social aim of egalitarianism. However, the regulatory attempts to bring artificial intelligence in line with liberal expectations have failed in many ways. Artificial intelligence is a paradigm of thought based on scientific discoveries, as well as a system of sociotechnical organization, as much as it is a specific set of technologies. In this section, we consider artificial intelligence in its own terms. We find that as a scientific field, it has proven and applied several ideas that are antithetical to liberalism, including control and optimization, bounded rationality, and the platform.

### 2.1 From subject to controller

Control engineering methods can be traced back to the end of the 19th century [13], and they developed through cybernetics—a post-WW2 interdisciplinary amalgam of emerging scientific theories and engineering practices [8, 35, 78]. Originally, control engineering concerned the creation of machines, like the homeostat or the heat-seeking missile, that would adjust their behavior dynamically to perform a function; for example, the missile rotates towards a target with speed that slows to zero as it nears the optimal orientation. Aligned with the new sciences of control came a new concept of the agent: the agent is a controller that behaves in an environment

to optimize its goals, often characterized by a von Neumann utility function.

Working with behavioral and cognitive scientists, cyberneticists soon began conceptualizing humans in these terms. This led to an idea of humans as “patterns of behavior” and symbolic information processors [35, 46], enabling the argument that “a uniform behavioristic analysis is applicable to both machines and living organisms, regardless of the complexity of the behavior.” Practitioners thus explored the utility of control, computation, and feedback for directing combined human-machine systems towards particular goals [78]. Cognitive scientists, always in close communication with AI research, have come to codify the laws of rational thought and action in the idiom of more basic statistical and computational theory [6, 24, 25, 50, 97]. Contemporary artificial intelligence is often framed as the implementation of these principles in machine systems designed to accomplish an expanding range of tasks. In variational Bayesian inference [40], for example, optimal beliefs are arrived at through iteratively tuning model parameters in an optimization procedure not unlike the targeting of a heat-seeking missile.

In this paradigm, the logical principles of individual human and artificial rationality are the same. The salient difference between the human, the organization, and artificial intelligence is therefore the ability to access data and process it efficiently. This view of the *agent as controller and controlled* flouts liberalism’s view of the individual as autonomous and private author of their own life. Instead, the individual cedes their special status and becomes one form of agent among many others.

### 2.2 From firm to organization

Whereas the firm for liberalism is a nexus of contracts, headed by the entrepreneur, made between rational individuals to reduce their transaction costs, artificial intelligence has transformed the firm into something else. Herbert Simon, the foundational AI theorist and economist, perhaps articulated this transformation best [85–87]. During World War 2, organizational decision-making became the subject of operations research optimization techniques like linear programming. Under these conditions, the rational choice assumptions of classical economic models, inspired by liberalism’s model of the rational individual subject, broke down. Rather, as organizations struggled with the computational intractability of the problems they faced, they discovered that people and machines were only boundedly rational. That meant problem-solving ability is limited by cognitive and sensory capacity, as well as environmental factors. Rationality thus became a complex, quantitative distinction, rather than a simple qualitative one as understood by liberalism.

With the understanding that humans are only boundedly rational has come the spectre of rationality that is greater-than-human. One version of this story is told by Bostrom [20], who argues that the world is threatened by the possibility of a superintelligence that could destroy human life as we know it. The argument goes that the first “general Artificial Intelligence,” or GAI, capable of thought approximating that of a human, would have a drive to expand its own cognitive power as an instrumental goal towards any other goals. Unless we can discover a way to perfectly align the superintelligence’s goals with humanity’s, its climb towards

omnipotence will endanger us all. While there are reasons to consider Bostrom's vision to be far-fetched [14], it can be interpreted as a compelling futurist parable about the modern corporation. Corporations, which unite many individuals and technical systems under centralized leadership, surely can outperform the individual at many cognitive tasks [86]. While originally a legal fiction designed to protect investors from liability, corporate personhood is now a legally recognized form of artificial agency that predates and informs questions about the nature of artificial intelligence today [99].

Economic theory has kept up with the changing nature of the firm. Alchian and Demsetz [4] proposed a revised theory of the firm in which the role of the manager is to be a 'meter' on internal work to prevent free-riding by workers. While Simon emphasized the goal-oriented nature of the firm, Alchian and Demsetz tilted the understanding of the firm towards that of a cybernetic unity. Today, liberalism's idea of the firm as an individual entrepreneur's contractual relations with other individuals is a quaint stage in the ontogenesis of the modern corporation, which typically has many absentee investors owning shares in a legal entity controlled by a managerial layer of organized people and information systems. The managerial layer is the system's Artificial Intelligence. More often than not, it is bound by fiduciary duty to the goal of maximizing profit for its shareholders [27]. To that end, it seeks greater control in, over, and as a market.

### 2.3 Smart markets

The widespread adoption of the Internet created booming potential for profit-seeking corporations. However, to make the most of these opportunities, businesses needed to monetize information services at scale. Artificial intelligence would soon be used to enable these new business models. According to the public market logic that underlies liberal theory, the price of a good should be equal to its marginal cost, which for information goods is understood to be near zero. How could a business profit from information goods on the Internet? One solution advanced by Shapiro and Varian [83] (that emerged alongside the expansion of intellectual property rights and their use for the extraction of rents on information economy goods) was effective price discrimination—charging customers individually the maximum price that they were willing to pay. Through individualized pricing, accomplished through the gathering of personalized purchase histories [2] and versioning of products, a firm can redirect what was consumer surplus into producer surplus and thereby recoup the fixed costs of producing the information good in the first place.

Hence, artificial intelligence has permeated and in many ways become the market. Whereas Coase saw the purpose of the firm as reducing transaction costs from market interactions, today some of the largest firms are themselves markets: platforms that offer a space for other businesses and consumers to operate with reduced transaction costs. Search engines, recommendation systems, and behavioral advertising all lubricate economic transactions. Both the market platforms and the digital businesses that operate on them constantly harvest information with which they can gain advantage over competitors and customers through massive scale experimentation [30].

But these markets are not public markets, wherein relevant information is available to all participants and private goods may be freely exchanged [31]. In this economy information services are rented to specific customers who do not have the opportunity to resell [47, 77]. While telecommunications infrastructure has been used in some places to increase price transparency and market efficiency [59], platforms operate in systems of profound information asymmetry, leading to unknown amounts of market failure [3, 94]. Some argue they are no longer even markets [18, 69, 100] but rather algorithmic simulations of markets designed to coordinate actors towards the goals of the market's corporate owner. This inverts the relationship between the public and private spheres. Under liberalism, individuals have private domains mediated by a public market; today, individual domains are made public, mediated by private platforms. This inversion changes the apparatus of economic distribution in liberalism that has been a mainstay of its legitimacy. Nothing guarantees equivalency between individuals in such a market; egalitarianism is undermined by the inequality between individuals and the platforms themselves.

## 3 THE PURPOSE OF SOCIAL SYSTEMS

Liberal law and artificial intelligence are working with fundamentally different materials. Liberalism imagines the autonomous rational individual, the public market, and a measure of equality guaranteed by the state. Artificial intelligence has given us unequally rational agents in hierarchies of corporate control that have subsumed market functions. If the units of liberal legal calculation cannot be made to function meaningfully in the artificial intelligence economy, the question thus becomes what might ethical or legal principles look like if they were to emerge from within artificial intelligence—especially in its organizational dimensions. The intellectual roots of artificial intelligence may lead to insights into new ethical thinking. One new terrain of political, ethical and regulatory possibility is the formation of new types of collectives and intermediaries. By intermediary, we do not mean government or regulatory intermediaries performing delegated tasks (i.e. as described by Abbott et al. [1]), or intermediate agents engaged to perform outsourced data processing without determining the purpose of that processing, or even entities certified to manage and deal with civic "open" data. Instead, here we use intermediaries to mean sociotechnical organizations positioned between individuals and other parts of the B2B "AI" economy, enabling the association of individuals around particular goals through the collective aggregation, mediation and operationalization of data flow.

In this section, we take a positive approach and introduce a new basis for artificial intelligence ethics that shares AI's understandings of agency as control, bounded rationality, and sociotechnical organization. The related fields of second order cybernetics, systems theory, and artificial life grew out of earlier cybernetics research which is now at the heart of contemporary machine learning systems. But while artificial intelligence devoted itself to the instrumental task of performing industrially relevant functions, this alternative intellectual branch has been taken up in legal theory and the computational understanding of biological systems. In particular, the idea of a social system (from Luhmann) provides a robust conceptualization of social form that is compatible with the

now ubiquitous conditions of technical mediation we see today. We propose, as a way forward, resolving the tension between liberalism and artificial intelligence by considering what sociotechnical organizations considered as a form of artificial life normatively imply. This ontological move leads concretely to the normative question: what shall be the purposes of our social systems? We see this as a key question for postliberal AI ethics that can be concretely mapped onto the design of data intermediaries.

### 3.1 From artificial intelligence to artificial life

Whereas cybernetics primarily studied the design of machines, second-order cybernetics began with the application of information and control theory to itself as a scientific field [68, 111] and more broadly to the understanding of biological life. The key ideas in second-order cybernetics, especially the version developed by Maturana and Varela, are that of autonomy, the ability of a system to maintain itself as a unity over time, and, as a special case of autonomy, autopoiesis, the ability of a system to create its own components and organization [105]. This fresh ontology, sometimes referred to as systems theory, does not take autonomy for granted, as both liberalism and artificial intelligence do, and rejects the idea that a system can be understood independently from its environment.

This line of reasoning has been quietly influential but never a dominant paradigm. Winograd and Flores [117] used this alternative paradigm in their influential critique of artificial intelligence when the latter was still in its symbolic stage. This inspired advances in the field of human-computer interaction (HCI), which centered on the situatedness of the user [33, 95, 116]. However, since their influential critique, artificial intelligence has reoriented around the technologies of the Internet and the smartphone, which have moved the locus of system design centripetally into the thick mesh of businesses bartering over decontextualized, aggregate consumer data.

Niklas Luhmann [66] built on Maturana and Varela's second order cybernetics to study social systems. While Luhmann is well-known in German social and legal scholarship [98], he is less well understood in Anglophone discourse (though he is used in e.g. Baxter, 1997, and Ho, 2013). Luhmann, building on Parsons, emphasizes the relative autonomy of different systems that are nevertheless structurally coupled. Two individuals, for example, each have their own psychic system, but they are structurally coupled with each other through a social system that is composed of their communications.

While engaged with philosophically heady and sociological concepts, this line of inquiry is nevertheless supported by a mathematized discipline. There are both classical mathematizations of autonomy [45, 107] and more recent attempts to rebuild the theory on information theoretic grounds [17]. Computational simulations are used to study artificial life [21] with models of agents whose cognition and identity are more intertwined with each other and their environment [106] than the von Neumann model of the agent familiar to artificial intelligence practitioners and emblematic of liberalism's atomistic conception of the individual. Second order cybernetics has been used in studies of software development and data modeling [39]), but not yet to more modern machine learning

applications, platform design, or theory of the firm. We see this as a promising research trajectory for artificial intelligence ethics with direct technical applications. Specifically, we consider the design of intermediaries: collections of individuals, joined by a technical system that engages them in a circuit of data flows, as a social system and form of artificial life.

### 3.2 Intermediary Purposes

As a historical contingency, today the dominant form of social organization has the profit function as its *raison d'être*. The individual can stand in only a limited number of relations towards such an organism: they can be shareholders, employees, customers, or, most recently, users. The interior relations between these parties, mediated today by artificial intelligence, are beyond the scope of liberal understanding and the normativity of text-based law and contracts [56]. Meanwhile, these organisms engage each other in business to business (B2B) relations, similarly mediated by artificial intelligence, in myriad complex ways involving calibrated, digital, probabilistic exchanges in which control, not commodities, are exchanged for money and entitlements.

Liberalism cannot describe the world with artificial intelligence, and so it cannot normatively guide it. Perhaps the design and legal facilitation of digital intermediaries can preserve some form of relative autonomy and equality for individuals. These designs should no longer be premised on the rationality of the individual. This rationality is dwarfed by those of social systems employing artificial intelligence. The problem is that those AI systems treat individuals as instruments towards the goal of profit. An ethical AI, in contrast, would be a social system that treats the individuals with which it is structurally coupled as ends, not means. More precisely, such a system would sustain itself, and its individual members as parts of itself, as a free way of life.

Data intermediaries, like privacy regulations, will fail if they are designed around mismatched liberal ideas. One of the earliest internet intermediaries, for instance, AllAdvantage, collected data and displayed personalized ads through a browser plugin. It remitted a fraction of the advertising value to its users. This commercial intermediary effectively invented behavioral advertising, but eventually collapsed because of its Ponzi-scheme-like financial model. Simply attempting to improve the individual's terms of trade over their data, imagined as a commodity, has never worked. Pistor [73] suggests a more radical solution of remitting value from data back to users as a form of equity by refashioning platforms into agents of user-owned data trusts. Through this mechanism data producers would become joint-venturers of the platform, giving them a claim on future profits as well as governance rights. This socializes the commercial platform, but leaves its commercial purposes otherwise consistent. Simultaneously occupying the roles of user and owner of a commercial platform apparatus may perversely incentivize individuals to exploit themselves to achieve maximum financial remuneration. The pursuit of profit for its own sake will perhaps never lead to an ethical AI.

Another purpose for intermediaries is the enforcement of data rights. One example is the UK App Drivers and Couriers union bringing a class action against Uber for breaching Article 22 of the GDPR in its use of automated decision-making systems to suspend

workers from the App [36]. While class actions leverage the economics of group representation, not-for-profits that bring these cases are still pursuing the enforcement of individual data subject rights as their primary purpose. Because there is little or no communicative action by class members with respect to legal action, it may not even be a social system per se. A more promising example is the Swiss health data coop MiData, which is positioned as a platform between providers of health data and entities interested in using it for different purposes. Users are given granular control over the use of their data in different projects. As members of the coop, users also have a stake in the development of the coop's governance, including how revenues paid into the coop are used. MiData is only one of several innovative organizational models for health data [61] an area of growing significance [108]. The MiData model is also being extended to other sectors, such as energy.

Sectoral data intermediaries may be more promising than those involved with context-free consumer internet, portable devices, or sensors in public. Quayside Civic Data Trust, intended for the Toronto Sidewalk Labs project, was a proposed data intermediary for the "smart city" urban data use case, for which it is not feasible to collect meaningful consent from data subjects [84]. It notionally involved a community controlled data review board with comprehensive internal governance mechanisms to develop policies and rules for data collection, release and processing. However, the "civic" purpose of the trust may have concealed another, the corporate risk strategy of Sidewalk Labs. Data's commercial value for advertising is time sensitive, but legal audit requirements require that it be maintained for some time beyond its commercial usefulness. The Civic Data Trust may thus have been in fact a Civic Data Dump for toxic data assets [12, 81]. The project never eventuated.

Many radical data intermediary design proposals focus on the decentralization of the technical architecture; these have met with varied success [70, 101]. Among the more recent proposals are Distributed Autonomous Organizations (DAO) that use blockchain technology to ensure that no one entity must be trusted to control the system. Using different cooperative arrangements, Distributed Cooperative Organizations (DisCOs) offer a theory of practice for recalibrating distributed technologies for facilitating associationist federations of cooperatives within alternative economies [102]. In the DisCo model, the social system tracks and remunerates different forms of work, including the labor necessary for community and cooperative reproduction (alongside generating commons-based and market-based forms of value). There are now a variety of modular technical applications and tools in the DisCo ecosystem for design, development, operation, automation, and reproduction of cooperatives.

In our view, an alternative model more aligned with second-order cybernetics principles would be something like an attention coop. Individuals that recognize the need to be part of a smart social system that reflects their shared purposes and priorities could join an intermediary in which members cooperatively develop their own search and recommendation priorities, controlling incoming data sources as well as outgoing data flows. Such a system might generate revenue if that were prioritized; following Pistor [73], that revenue could be distributed in pro rata shares to the system's members. But profit need not be the only purpose. Rather the purpose of the system would be self-control and self-governance. This narrow

function could be a prerequisite to building intermediaries that calibrated such systems towards civic priorities and public service, broadly construed, as suggested by some commentators [43, 80]. Perhaps such a media system would avoid the pathologies associated with social media optimized for engagement for advertising purposes.

In general, we see the movement from liberalism towards second-order cybernetics as opening a wider range of social systems as potential solutions to the challenges of AI law and ethics. Rather than attempt to remit data's value to individuals through a property right, social systems can leverage their market power and collectively bargain for better terms of interaction with the digital economy. Rather than assert rights through class action lawsuits, social systems can engage in collective action to steer political systems. Rather than limit themselves to compliance with existing sectoral regulations, social systems can create new contexts with new social purposes. While intermediaries are often launched by entrepreneurs and activists, the viability of new intermediary forms may well depend on the legal environment. Our inquiry leads to the recommendation that AI ethicists look to the field of organization law [73] and explore it for new possibilities.

We have arrived at a principle: an artificial intelligence will be as ethical as the purposes of the social system that operates it. AI ethics implicates many broader political and sociological questions about the organization of society, but demands attention to the scientific advances at the heart of artificial intelligence practice. This science has revealed the limitations of atomistic individuality and drawn our attention to collective interests. If data protection law and its analogues are premised on individual autonomy and control for the protection of dignity, personhood, and self-presentation, what purposes might constitute meaningful collective autonomy in the data economy? Is autonomy itself such a purpose?

## ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees for their valuable comments and helpful suggestions. This work has been immeasurably improved by the kind and thoughtful attention of the research groups at Cornell Tech's Digital Life Initiative and NYU School of Law's Information Law Institute and Privacy Research Group, with special thanks to Katherine Strandburg, Helen Nissenbaum, Lee McGuigan, Salomé Viljoen, Elettra Bietti, Ira Rubinstein, Ari Waldman, and Thomas Streinz among others. An earlier version was presented at the 2019 European Privacy Law Scholars Conference, and another prepared for the 2020 Ethics of Data Science Conference, and we thank the reviewers for their comments. A version of this work was presented at a seminar arranged by the Centre for AI and Digital Ethics at Melbourne Law, School, University of Melbourne, and the Humanizing Machine Intelligence group at the Australian National University, and we thank Jeannie Paterson and Seth Lazar for helpful comments.

## REFERENCES

- [1] Kenneth W Abbott, David Levi-Faur, and Duncan Snidal. 2017. Theorizing regulatory intermediaries: The RIT model. *The ANNALS of the American Academy of Political and Social Science* 670, 1 (2017), 14–35.
- [2] Alessandro Acquisti and Hal R Varian. 2005. Conditioning prices on purchase history. *Marketing Science* 24, 3 (2005), 367–381.



- [3] George A Akerlof. 1978. The market for “lemons”: Quality uncertainty and the market mechanism. In *Uncertainty in Economics*. Elsevier, 235–251.
- [4] Armen A Alchian and Harold Demsetz. 1972. Production, information costs, and economic organization. *The American Economic Review* 62, 5 (1972), 777–795.
- [5] Elizabeth Anderson. 2009. Toward a non-ideal, relational methodology for political philosophy: Comments on Schwartzman’s “Challenging Liberalism”. *Hypatia* 24, 4 (2009), 130–145.
- [6] John R Anderson. 1991. Is human cognition adaptive? *Behavioral and Brain Sciences* 14, 3 (1991), 471–485.
- [7] Imanol Arrieta-Ibarra, Leonard Goff, Diego Jiménez-Hernández, Jaron Lanier, and E Glen Weyl. 2018. Should we treat data as labor? Moving beyond “free”. In *American Economic Association Papers and Proceedings*, Vol. 108. 38–42.
- [8] W Ross Ashby. 1940. Adaptiveness and equilibrium. *Journal of Mental Science* 86, 362 (1940), 478–483.
- [9] Solon Barocas and Helen Nissenbaum. 2009. On notice: The trouble with notice and consent. In *Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information*.
- [10] Lemi Baruh and Mihaela Popescu. 2017. Big data analytics and the limits of privacy self-management. *New Media & Society* 19, 4 (2017), 579–596.
- [11] Hugh Baxter. 1997. Autopoiesis and the relative autonomy of law. *Cardozo L. Rev.* 19 (1997), 1987.
- [12] Omri Ben-Shahar. 2019. Data pollution. *Journal of Legal Analysis* 11 (2019), 104–159.
- [13] James Beniger. 2009. *The Control Revolution: Technological and Economic Origins of the Information Society*. Harvard university press.
- [14] Sebastian Benthall. 2017. Don’t fear the reaper: refuting Bostrom’s superintelligence argument. *arXiv preprint arXiv:1702.08495* (2017).
- [15] Sebastian Benthall and Jake Goldenfein. 2020. Essential infrastructures. *Phenomenal World* (2020). <https://phenomenalworld.org/analysis/essential-infrastructures>
- [16] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2018. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research* (2018), 0049124118782533.
- [17] Nils Bertschinger, Eckehard Olbrich, Nihat Ay, and Jürgen Jost. 2008. Autonomy: An information theoretic perspective. *Biosystems* 91, 2 (2008), 331–345.
- [18] Kean Birch. 2020. Automated Neoliberalism? The Digital Organisation of Markets in Technoscientific Capitalism. *New Formations* 100, 100–101 (2020), 10–27.
- [19] Kean Birch and Fabian Muniesa. 2020. *Assetization: Turning Things into Assets in Technoscientific Capitalism*. MIT Press.
- [20] Nick Bostrom. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- [21] Paul Bourgin and Francisco J Varela. 1992. Towards a practice of autonomous systems. In *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, Vol. 1. Cambridge, MA: MIT Press.
- [22] William W Bratton. 1988. Nexus of contracts corporation: A critical appraisal. *Cornell Law Review* 74 (1988), 407.
- [23] Jebediah Britton-Purdy, David Singh Grewal, Amy Kapczynski, and K. Sabeel Rahman. 2020. Building a law-and-political-economy framework: Beyond the twentieth-century synthesis. *Yale Law Journal* 129, 6 (2020), 1784–1835.
- [24] Nick Chater and Mike Oaksford. 1999. Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences* 3, 2 (1999), 57–65.
- [25] Nick Chater and Paul Vitányi. 2003. Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences* 7, 1 (2003), 19–22.
- [26] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* 5, 2 (2017), 153–163.
- [27] David Ciepley. 2019. Wayward leviathans: How America’s corporations lost their public purpose. *The Hedgehog Review* 21, 1 (2019), 68–85.
- [28] Marika Cifor, Patricia Garcia, TL Cowan, Jasmine Rault, Tonia Sutherland, A Chan, Jennifer Rode, Anna Lauren Hoffmann, Niloufar Salehi, and Lisa Nakamura. 2019. *Feminist Data Manifest-No.* (2019). <https://www.manifestno.com/>
- [29] Ronald Harry Coase. 1995. The nature of the firm. In *Essential Readings in Economics*. Springer, 37–54.
- [30] Julie E Cohen. 2017. Law for the platform economy. *University of California Davis Law Review* 51 (2017), 133.
- [31] Julie E Cohen. 2017. Surveillance vs. privacy: effects and implications. In *Cambridge Handbook of Surveillance Law*, David Gray and Stephen E Henderson (Eds.). Cambridge University Press: New York, 455–69.
- [32] Julie E Cohen. 2019. *Between Truth and Power: The Legal Constructions of Informational Capitalism*. Oxford University Press, USA.
- [33] Paul Dourish. 2004. *Where the Action Is: The Foundations of Embodied Interaction*. MIT press.
- [34] Bernard Edelman. 1979. *Ownership of the Image: Elements for a Marxist Theory of Law*. Routledge & Kegan Paul London.
- [35] Paul N Edwards. 1997. *The Closed World: Computers and the Politics of Discourse in Cold War America*. MIT Press.
- [36] Ekker. 2020. APPLICATION pursuant to Art. 22(1) GDPR and Art. 35 UAVG - District Court of Amsterdam. <https://ekker.legal/wp-content/uploads/2020/10/Court-request-Uber-account-deactivation-unofficial-translation.pdf>
- [37] Barbara J Evans. 2017. Power to the people: data citizens in the age of precision medicine. *Vanderbilt Journal of Entertainment and Technology Law* 19, 2 (2017), 243.
- [38] John Finnis. 1993. Liberalism and natural law theory. *Mercer Law Review* 45 (1993), 687.
- [39] Christiane Floyd, Heinz Züllighoven, Reinhard Budde, and Reinhard Keil-Slawik. 1992. *Software Development and Reality Construction*. Berlin: Springer.
- [40] Charles W Fox and Stephen J Roberts. 2012. A tutorial on variational Bayesian inference. *Artificial Intelligence Review* 38, 2 (2012), 85–95.
- [41] Samuel Freeman. 2001. Illiberal libertarians: Why libertarianism is not a liberal view. *Philosophy & Public Affairs* 30, 2 (2001), 105–151.
- [42] Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. 2016. On the (im) possibility of fairness. *arXiv preprint arXiv:1609.07236* (2016).
- [43] Christian Fuchs. 2014. Social media and the public sphere. *tripleC: Communication, Capitalism & Critique. Open Access Journal for a Global Sustainable Information Society* 12, 1 (2014), 57–101.
- [44] Oscar H Gandy Jr. 1993. *The Panoptic Sort: A Political Economy of Personal Information (Critical Studies in Communication and in the Cultural Industries)*. Westview Press.
- [45] Joseph A Goguen and Francisco J Varela. 1979. Systems and distinctions; duality and complementarity. *International Journal of General System* 5, 1 (1979), 31–43.
- [46] Jake Goldenfein. 2019. *Monitoring Laws: Profiling and Identity in the World State*. Cambridge University Press: New York.
- [47] Jake Goldenfein. 2019. RAM-Chip. In *A History of Intellectual Property in 50 Objects*, Claudy Op Den Kamp and Dan Hunter (Eds.). Vol. 50. Cambridge University Press: New York.
- [48] Mary L Gray and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan Books.
- [49] Ben Green and Salomé Viljoen. 2020. Algorithmic realism: expanding the boundaries of algorithmic thought. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 19–31.
- [50] Thomas Griffiths, Charles Kemp, and Joshua B Tenenbaum. 2008. Bayesian models of cognition. In *Cambridge Handbook of Computational Cognitive Modeling*. Carnegie Mellon University.
- [51] Robert L Hale. 1923. Coercion and distribution in a supposedly non-coercive state. *Political Science Quarterly* 38, 3 (1923), 470–494.
- [52] Moritz Hardt, Eric Price, and Nathan Srebro. 2016. Equality of opportunity in supervised learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*. 3323–3331.
- [53] Friedrich A Hayek. 1973. *Rules and Order. Vol. 1 of Law, Legislation, and Liberty*. Chicago: University of Chicago Press.
- [54] Yoan Hermstruwer. 2017. Contracting around privacy: the (behavioral) law and economics of consent and big data. *JIPITEC* 8 (2017), 9.
- [55] Mireille Hildebrandt. 2008. Defining profiling: a new type of knowledge? In *Profiling the European citizen*. Springer, 17–45.
- [56] Mireille Hildebrandt. 2021. The adaptive nature of text-driven law. *Journal of Cross-disciplinary Research in Computational Law* 1, 1 (2021).
- [57] Christina S Ho. 2013. In defense of circular reasoning: The Affordable Care Act and the resilience of law and self-reference. *William and Mary Policy Review* 5 (2013), 1.
- [58] Chris Jay Hoofnagle and Jan Whittington. 2013. Free: accounting for the costs of the internet’s most popular price. *UCLA L. Rev.* 61 (2013), 606.
- [59] Robert Jensen. 2007. The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector. *The Quarterly Journal of Economics* 122, 3 (2007), 879–924.
- [60] Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2016. Fair algorithms for infinite and contextual bandits. *arXiv preprint arXiv:1610.09559* (2016).
- [61] Timothy Kariotis, Mad Price Ball, Bastian Greshake Tzovaras, Simon Dennis, Tony Sahama, Carolyn Johnston, Helen Almond, and Ann Borda. 2020. Emerging health data platforms: From individual control to collective data governance. *Data & Policy* 2 (2020).
- [62] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807* (2016).
- [63] Donald P Kommers. 1994. The federal constitutional court in the German political system. *Comparative Political Studies* 26, 4 (1994), 470–491.
- [64] Jaron Lanier. 2014. *Who Owns the Future?* Simon and Schuster.
- [65] Kenneth C Laudon. 1996. Markets and privacy. *Commun. ACM* 39, 9 (1996), 92–104.
- [66] Niklas Luhmann. 1995. *Social Systems*. Stanford University Press.
- [67] Alfred Marshall. 2009. *Principles of Economics: Unabridged Eighth Edition*. Cosimo, Inc.
- [68] Margaret Mead. 1968. The Cybernetics of Cybernetics. In *Purposive Systems*, Larry J. Peterson Heinz von Foerster, John D. White and John K. Russell (Eds.). Spartan Books, New York, NY, USA. 1–11.
- [69] Philip Mirowski and Edward Nik-Khah. 2017. *The Knowledge We Have Lost in Information: The History of Information in Modern Economics*. Oxford University

- Press.
- [70] Arvind Narayanan, Vincent Toubiana, Solon Barocas, Helen Nissenbaum, and Dan Boneh. 2012. A critical look at decentralized personal data architectures. *Workshop on Data Usage Management on the Web (DUMW)* (2012).
- [71] Helen Nissenbaum. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press.
- [72] Evgeny Pashukanis. 2017. *The General Theory of Law and Marxism*. Routledge.
- [73] Katharina Pistor. 2020. Rule by data: The end of markets? *Law & Contemporary Problems* 83 (2020), 101.
- [74] Julia Powles and Helen Nissenbaum. 2018. The seductive diversion of ‘solving’ bias in artificial intelligence. <https://medium.com/s/story/these-deductive-diversion-of-solving-bias-in-artificial-intelligence-890df5e5ef53>
- [75] K Sabeel Rahman. 2018. Regulating informational infrastructure: Internet platforms as the new public utilities. *Georgetown Law and Technology Review* 2 (2018), 2.
- [76] John Rawls. 2009. *A Theory of Justice*. Harvard University Press.
- [77] Jeremy Rifkin. 2001. *The Age of Access: The New Culture of Hypercapitalism*. Penguin.
- [78] Arturo Rosenblueth, Norbert Wiener, and Julian Bigelow. 1943. Behavior, purpose and teleology. *Philosophy of Science* 10, 1 (1943), 18–24.
- [79] Pamela Samuelson. 2000. Privacy as intellectual property? *Stanford Law Review* (2000), 1125–1173.
- [80] Dan Schiller. 2020. Reconstructing public utility networks: A program for action. *International Journal of Communication* 14 (2020), 12.
- [81] Bruce Schneier. 2016. Data is a Toxic Asset. [https://www.schneier.com/blog/archives/2016/03/data\\_is\\_a\\_toxic.html](https://www.schneier.com/blog/archives/2016/03/data_is_a_toxic.html)
- [82] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 59–68.
- [83] Carl Shapiro and Hal R Varian. 1998. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business Press.
- [84] Sidewalk Labs. 2018. Digital Governance Proposals for DSAP Consultation. [https://waterfrontoronto.ca/nbe/wcm/connect/waterfront/41979265-8044-442a-9351-e28ef6c76d70/18.10.15\\_SWT\\_Draft+Proposals+Regarding+Data+Use+and+Governance.pdf](https://waterfrontoronto.ca/nbe/wcm/connect/waterfront/41979265-8044-442a-9351-e28ef6c76d70/18.10.15_SWT_Draft+Proposals+Regarding+Data+Use+and+Governance.pdf)
- [85] Herbert A Simon. 1957. *Models of Man: Social and Rational*. Wiley.
- [86] Herbert A Simon. 1972. Theories of bounded rationality. *Decision and Organization* 1, 1 (1972), 161–176.
- [87] Herbert A Simon. 2019. *The Sciences of the Artificial*. MIT press.
- [88] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2020. Participation is not a design fix for machine learning. International Conference on Machine Learning Workshop. In *Participatory Approaches to Machine Learning*.
- [89] Quinn Slobodian. 2020. *Globalists: The End of Empire and the Birth of Neoliberalism*. Harvard University Press.
- [90] Sarah Spiekermann, Alessandro Acquisti, Rainer Böhme, and Kai-Lung Hui. 2015. The challenges of personal data markets and privacy. *Electronic Markets* 25, 2 (2015), 161–167.
- [91] Dina Srinivasan. 2019. The antitrust case against Facebook: A monopolist’s journey towards pervasive surveillance in spite of consumers’ preference for privacy. *Berkeley Business Law Journal* 16 (2019), 39.
- [92] Dina Srinivasan. 2020. Why Google dominates advertising markets: Competition policy should lean on the principles of financial market regulation. *Stanford Technology Law Review* 24 (2020), 55–175.
- [93] Luke Stark. 2019. Facial recognition is the plutonium of AI. *XRDS: Crossroads, The ACM Magazine for Students* 25, 3 (2019), 50–55.
- [94] Katherine J Strandburg. 2013. Free fall: The online market’s consumer preference disconnect. *University of Chicago Legal Forum* (2013), 95.
- [95] Lucy Suchman. 1993. Do categories have politics? The language/action perspective reconsidered. In *Proceedings of the Third European Conference on Computer-Supported Cooperative Work 13–17 September 1993, Milan, Italy ECSCW’93*. Springer, 1–14.
- [96] Brian Z Tamanaha. 2004. *On the rule of law: History, politics, theory*. Cambridge University Press.
- [97] Joshua B Tenenbaum, Charles Kemp, Thomas L Griffiths, and Noah D Goodman. 2011. How to grow a mind: Statistics, structure, and abstraction. *Science* 331, 6022 (2011), 1279–1285.
- [98] Gunther Teubner. 1993. *Law as an Autopoietic System*. Oxford/Cambridge, Blackwell Publishers.
- [99] Gunther Teubner. 2006. Rights of non-humans? Electronic agents and animals as new actors in politics and law. *Journal of Law and Society* 33, 4 (2006), 497–521.
- [100] Julia Tomassetti. 2016. Does Uber redefine the firm: the postindustrial corporation and advanced information technology. *Hofstra Labor & Employment Law Journal* 34 (2016), 1.
- [101] Vincent Toubiana, Arvind Narayanan, Dan Boneh, Helen Nissenbaum, and Solon Barocas. 2010. Adnostic: Privacy preserving targeted advertising. In *Proceedings Network and Distributed System Symposium*.
- [102] Stacco Troncoso and Ann Marie Utratel. 2019. If I Only Had a Heart: A DisCO Manifesto. *The Transnational Institute* 14 (2019).
- [103] D Tucker. 2012. *Essay on Liberalism: Looking Left and Right*. Vol. 51. Springer Science & Business Media.
- [104] Alan M Turing. 2009. Computing machinery and intelligence. In *Parsing the Turing Test*. Springer, 23–65.
- [105] Francisco G Varela, Humberto R Maturana, and Ricardo Uribe. 1974. Autopoiesis: the organization of living systems, its characterization and a model. *Biosystems* 5, 4 (1974), 187–196.
- [106] Francisco J Varela. 1997. Patterns of life: Intertwining identity and cognition. *Brain and Cognition* 34, 1 (1997), 72–87.
- [107] Francisco J Varela and Joseph A Goguen. 1978. The arithmetic of closure. *Cybernetics and System* 8, 3-4 (1978), 291–324.
- [108] Paraskevas Vezyridis and Stephen Timmons. 2021. E-Infrastructures and the divergent assetization of public health data: Expectations, uncertainties, and asymmetries. *Social Studies of Science* 0, 0 (2021). <https://doi.org/10.1177/0306312721989818>
- [109] Salomé Viljoen. 2020. Data as Property? *Phenomenal World* (Oct 2020). <https://phenomenalworld.org/analysis/data-as-property>
- [110] Salomé Viljoen. 2021 (forthcoming). Democratic Data: A Relational Theory for Data Governance. *Yale Law Journal* (2021 (forthcoming)).
- [111] Heinz Von Foerster. 2003. Cybernetics of cybernetics. In *Understanding understanding*. Springer, 283–286.
- [112] Robin West. 1986. Law, Rights, and Other Totemic Illusions: Legal Liberalism and Freud’s Theory of the Rule of Law. *University of Pennsylvania Law Review* 134, 4 (1986), 817–882.
- [113] Alan F Westin. 1968. Privacy and freedom. *Washington and Lee Law Review* 25, 1 (1968), 166.
- [114] James Q Whitman. 2003. The two western cultures of privacy: Dignity versus liberty. *Yale Law Journal* 113 (2003), 1151.
- [115] Oliver E Williamson. 1989. Transaction cost economics. *Handbook of Industrial Organization* 1 (1989), 135–182.
- [116] Terry Winograd. 1993. Categories, disciplines, and social coordination. *Computer Supported Cooperative Work (CSCW)* 2, 3 (1993), 191–197.
- [117] Terry Winograd and Fernando Flores. 1986. *Understanding Computers and Cognition: A New Foundation for Design*. Intellect Books.
- [118] Samuel Yeom and Michael Carl Tschantz. 2021. Avoiding Disparity Amplification under Different Worldviews. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 273–283.
- [119] M Zey. 2001. Rational Choice and Organization Theory. *International Encyclopedia of the Social & Behavioral Sciences* (2001).