The Moral Economy of High-Tech Modernism

Henry Farrell & Marion Fourcade

While people in and around the tech industry debate whether algorithms are political at all, social scientists take the politics as a given, asking instead how this politics unfolds: how algorithms concretely govern. What we call "high-tech modernism" – the application of machine learning algorithms to organize our social, economic, and political life – has a dual logic. On the one hand, like traditional bureaucracy, it is an engine of classification, even if it categorizes people and things very differently. On the other, like the market, it provides a means of self-adjusting allocation, though its feedback loops work differently from the price system. Perhaps the most important consequence of high-tech modernism for the contemporary moral political economy is how it weaves hierarchy and data-gathering into the warp and woof of everyday life, replacing visible feedback loops with invisible ones, and suggesting that highly mediated outcomes are in fact the unmediated expression of people's own true wishes.

Algorithms – especially machine learning algorithms – have become major social institutions. To paraphrase anthropologist Mary Douglas, algorithms "do the classifying."[1] They assemble and they sort – people, events, things. They distribute material opportunities and social prestige. But do they, like all artifacts, have a particular politics?[2] Technologists defend themselves against the very notion, but a lively literature in philosophy, computer science, and law belies this naive view. Arcane technical debates rage around the translation of concepts such as fairness and democracy into code. For some, it is a matter of legal exposure. For others, it is about designing regulatory rules and verifying compliance. For a third group, it is about crafting hopeful political futures.[3]

The questions from the social sciences are often different: How do algorithms concretely govern? How do they compare to other modes of governance, like bureaucracy or the market? How does their mediation shape moral intuitions, cultural representations, and political action? In other words, the social sciences worry not only about specific algorithmic outcomes, but also about the broad, society-wide consequences of the deployment of algorithmic regimes – systems of decision-making that rely heavily on computational processes running on large

databases. These consequences are not easy to study or apprehend. This is not just because, like bureaucracies, algorithms are simultaneously rule-bound and secretive. Nor is it because, like markets, they are simultaneously empowering and manipulative. It is because they are a bit of both. Algorithms extend *both* the logic of hierarchy *and* the logic of competition. They are machines for making categories and applying them, much like traditional bureaucracy. And they are self-adjusting allocative machines, much like canonical markets.

Understanding this helps highlight both similarities and differences between the historical regime that political scientist James Scott calls "high modernism" and what we dub *high-tech modernism*.[4] We show that bureaucracy, the typical high modernist institution, and machine learning algorithms, the quintessential high-tech modernist one, share common roots as technologies of hierarchical classification and intervention. But whereas bureaucracy reinforces human sameness and tends toward large, monopolistic (and often state-based) organizations, algorithms encourage human competition, in a process spearheaded by large, near-monopolistic (and often market-based) organizations. High-tech modernism and high modernism are born from the same impulse to exert control, but are articulated in fundamentally different ways, with quite different consequences for the construction of the social and economic order. The contradictions between these two moral economies, and their supporting institutions, generate many of the key struggles of our times.

Both bureaucracy and computation enable an important form of social power: the power to classify.[5] Bureaucracy deploys filing cabinets and memorandums to organize the world and make it "legible," in Scott's terminology. Legibility is, in the first instance, a matter of classification. Scott explains how "high modernist" bureaucracies crafted categories and standardized processes, turning rich but ambiguous social relationships into thin but tractable information. The bureaucratic capacity to categorize, organize, and exploit this information revolutionized the state's ability to get things done. It also led the state to reorder society in ways that reflected its categorizations and acted them out. Social, political, and even physical geographies were simplified to make them legible to public officials. Surnames were imposed to tax individuals; the streets of Paris were redesigned to facilitate control.

Yet high modernism was not just about the state. Markets, too, were standardized, as concrete goods like grain, lumber, and meat were converted into abstract qualities to be traded at scale.[6] The power to categorize made and shaped markets, allowing grain buyers, for example, to create categories that advantaged them at the expense of the farmers they bought from. Businesses created their own bureaucracies to order the world, deciding who could participate in markets and how goods ought to be categorized.

We use the term high-tech modernism to refer to the body of classifying technologies based on quantitative techniques and digitized information that partly displaces, and partly is layered over, the analog processes used by high modernist organizations. Computational algorithms – especially machine learning algorithms – perform similar functions to the bureaucratic technologies that Scott describes. Both supervised machine learning (which classifies data using a labeled training set) and unsupervised machine learning (which organizes data into self-discovered clusters) make it easier to categorize unstructured data at scale. But unlike their paper-pushing predecessors in bureaucratic institutions, the humans of high-tech modernism disappear behind an algorithmic curtain. The workings of algorithms are much less visible, even though they penetrate deeper into the social fabric than the workings of bureaucracies. The development of smart environments and the Internet of Things has made the collection and processing of information about people too comprehensive, minutely geared, inescapable, and fast-growing for considered consent and resistance.

In a basic sense, machine learning does not strip away nearly as much information as traditional high modernism. It potentially fits people into categories ("classifiers") that are narrower – even bespoke. The movie streaming platform Netflix will slot you into one of its two thousand–plus "microcommunities" and match you to a subset of its thousands of subgenres. Your movie choices alter your position in this scheme and might in principle even alter the classificatory grid itself, creating a new category of viewer reflecting your idiosyncratic viewing practices.

Many of the crude, broad categories of nineteenth-century bureaucracies have been replaced by new, multidimensional classifications, powered by machine learning, that are often hard for human minds to grasp.[7] People can find themselves grouped around particular behaviors or experiences, sometimes ephemeral, such as followers of a particular YouTuber, subprime borrowers, or fans of action movies with strong female characters. Unlike clunky high modernist categories, high-tech modernist ones can be *emergent* and technically *dynamic*, adapting to new behaviors and information as they come in. They incorporate tacit information in ways that are sometimes spookily right, and sometimes disturbing and misguided: music-producing algorithms that imitate a particular artist's style, language models that mimic social context, or empathic AI that supposedly grasps one's state of mind.[8] Generative AI technologies can take a prompt and generate an original picture, video, poem, or essay that seems to casual observers as though it were produced by a human being.

Taken together, these changes foster a new politics. Traditional high modernism did not just rely on standard issue bureaucrats. It empowered a wide variety of experts to make decisions in the area of their particular specialist knowledge and authority. Now, many of these experts are embattled, as their authority is nibbled

away by algorithms whose advocates claim are more accurate, more reliable, and less partial than their human predecessors.

One key difference between the moral economies of high modernism and high-tech modernism involves feedback. It is tempting to see high modernism as something imposed entirely from above. However, in his earlier book *Weapons of the Weak*, Scott suggests that those at the receiving end of categorical violence are not passive and powerless.[9] They can sometimes throw sand into the gears of the great machinery.

As philosopher Ian Hacking explains, certain kinds of classifications – typically those applying to human or social collectives – are "interactive" in that

> when known by people or those around them, and put to work in institutions, [they] change the ways in which individuals experience themselves – and may even lead people to evolve their feelings and behavior in part because they are so classified.[10]

People, in short, have agency. They are not submissive dupes of the categories that objectify them. They may respond to being put in a box by conforming to or growing into those descriptions. Or they may contest the definition of the category, its boundaries, or their assignment to it.[11] This creates a feedback loop in which the authors of classifications (state officials, market actors, experts from the professions) may adjust the categories in response. Human society, then, is forever being destructured and restructured by the continuous interactions between classifying institutions and the people and groups they sort.

But conscious agency is only possible when people know about the classifications: the politics of systems in which classifications are visible to the public, and hence potentially actionable, will differ from the politics of systems in which they are not.

So how does the change from high modernism to high-tech modernism affect people's relationships with their classifications? At its worst, high modernism stripped out tacit knowledge, ignored public wishes and public complaints, and dislocated messy lived communities with sweeping reforms and grand categorizations, making people *more visible* and hence more readily acted on. The problem was not that the public did not notice the failures, but that their views were largely ignored. Authoritarian regimes constricted the range of ways in which people could respond to their classification: anything more than passive resistance was liable to meet brutal countermeasures. Democratic regimes were, at least theoretically, more open to feedback, but often ignored it when it was inconvenient and especially when it came from marginalized groups.

The pathologies of computational algorithms are often more subtle. The shift to high-tech modernism allows the means of ensuring legibility to fade into the background of the ordinary patterns of our life. Information gathering is woven

into the warp and woof of our existence, as entities gather ever finer data from our phones, computers, doorbell cameras, purchases, and cars. There is no need for a new Haussmann to transform cramped alleyways into open boulevards, exposing citizens to view.[12] Urban architectures of visibility have been rendered nearly redundant by the invisible torrents of data that move through the air, conveying information about our movements, our tastes, and our actions to be sieved through racks of servers in anonymous, chilled industrial buildings.

The feedback loops of high-tech modernism are also structurally different. Some kinds of human feedback are now much less common. Digital classification systems may group people in ways that are not always socially comprehensible (in contrast to traditional categories such as female, married, Irish, or Christian). *Human* feedback, therefore, typically requires the mediation of specialists with significant computing expertise, but even they are often mystified by the operation of systems they have themselves designed.[13]

The political and social mechanisms through which people previously responded, actively and knowingly, to their categorization – by affirming, disagreeing with, or subverting it – have been replaced by closed loops in which algorithms assign people unwittingly to categories, assess their responses to cues, and continually update and reclassify them. The classifications produced by machine learning are cybernetic, in mathematician Norbert Wiener's original sense of the word. That is, they are self-correcting: categories are automatically and dynamically adjusted in light of the reactions that they produce.

The changing politics of credit in the United States helps illuminate these differences. Until the 1970s, broad demographic characteristics such as gender or race – or high modernist proxies such as marital status or the redlining of poor, primarily Black neighborhoods – were routinely used to determine a person's creditworthiness. It is only when categorical discrimination was explicitly forbidden that new actuarial techniques, aimed at precisely scoring the "riskiness" of specific individuals, started to flourish in the domain of credit.[14]

This did not just change how lenders "saw" individuals and groups, but also how individuals and groups thought about themselves and the politics that were open to them.[15] Redlining was overt racial prejudice, visible to anyone who bothered looking at a map. But credit scoring turned lending risk evaluation into a quantitative, individualized, and abstract process. Contesting the resulting classifications or acting collectively against them became harder. Later, the deployment of machine learning – which uses even weaker signals to make its judgments, like using one's phone's average battery level to determine their likelihood to repay their loan – made the process of measuring creditworthiness even more opaque and difficult to respond to.[16]

Predictive scores that rely on behavioral measures eschew blatant racial *discrimination*. But it would be a mistake to think that they eliminate racial *disparities* –

they just make them harder to see, sometimes allowing them to ramify further.[17] This is why the political struggle against algorithms has emphasized historical biases embedded in training data sets and the inherent unfairness and poor performance of nontransparent, automated decision-making. The European Commission has proposed to regulate the use of "high risk" algorithms that endanger fundamental rights, subjecting them to frequent human review.[18] This would include the use of algorithms for public benefit eligibility, credit scoring, law enforcement, immigration control, employment, and more. Finally, traditional high modernist professionals – including judges, journalists, and law enforcement officers – have also pushed back against the use of algorithms in their work, treating them as irrelevant, inefficient, or a status threat.[19]

The moral economy of high-tech modernism is market-driven, both practically and ideologically. Many algorithm-based start-ups want to expand market share rapidly and aggressively. Once revenues exceed fixed costs, the additional cost of adding a new user is comparatively tiny. Platform companies like Facebook or YouTube can serve billions of customers with tens of thousands of employees. Machine learning algorithms can gather data about users and dynamically provide and adjust flows of content, while auction and matching algorithms can maintain dynamic markets for advertisers who want access to customers with specific demographic characteristics.

Algorithms institutionalize competition between units (whether people, organizations, or ideas) by fostering a market-based vision of fairness.[20] The threat of being automated away looms large for all workers. Algorithmic technologies can also be implemented to hire and fire, to predict performance, influence, and riskiness, or to surveil, discipline, and arrest. They do so by rank-ordering according to their own particular versions of merit.[21] It is as though anyone who applies themselves can do well, and social structure and existing power allocations did not matter. (The irony is that while high-tech modernist firms are happy to turn the market screw on everyone else, they strive to establish monopoly for themselves).[22]

Just like the behavior of individuals, the distribution of knowledge must be subjected to the market test. High-tech modernism claims to represent popular judgment against the snobbishness of elites. Remember that Scott identifies high modernism as inherently antidemocratic because it enforces categories and objectives decided on by elites who "know better."[23] High-tech modernism, by contrast, systematically undermines elite judgment, fueling a crisis of expertise.[24] Algorithms purport to read X-rays better than radiologists, predict purchases better than market researchers, understand people's sexuality better than they themselves do, and produce new text or code better than many professional writers and engineers. Meanwhile, they elevate a kind of bottom-up wisdom. The network leaves it up to the crowd to judge what is worth knowing, generating collective

sentiments through likes, clicks, and comments. Viral trends and online multitudes provide a kind of pseudodemocratic, if extremely volatile, *vox populi*.

The absence of visible hierarchy legitimates high-tech modernism's claim that clouds and crowds best represent people's wishes. Its new elites echo early libertarian arguments about cyberspace, and quasi-Hayekian defenses of the market, facially justifying the notion that search engines and other algorithms are disinterested means of processing the internet's naturally dispersed stock of knowledge.[25] They flatter high-tech modernism as defending the liberties of the individual, freed from physical and social bonds, against traditional status hierarchies. The abundant data that people "freely" upload or leave behind as they roam cyberspace become "an unqualified good," fostering beneficial competition for everyone and everything.[26]

The awkward fact is that hierarchy has not disappeared. It has only become less visible. Platform companies' business priorities determine the algorithms that are employed, as well as their "objective functions," the weighted goals that they are supposed to maximize on. Social media corporations employ algorithms that maximize "engagement," keeping consumers scrolling through feeds or watching video clips so that they keep seeing paid content that may itself be misleading. Amazon, in contrast, cares more about getting people to buy things, and, according to legal scholar and Federal Trade Commission Chair Lina Khan, uses its detailed transaction information and ability to rank search outcomes to fortify its market dominance.[27] Platform companies dislike even tweaking their algorithms in response to regulators' demands for fear that it might create a precedent for further interventions that would conflict with their business model.

As search engines have transformed from general-purpose technology to personal digital assistants, they have elevated searching the web and forming an opinion "for oneself" into a normative principle. People think of search engines as oracles, but as sociologist Francesca Tripodi and others have shown, they work more like distorting mirrors that variously confirm, exacerbate, or take advantage of people's priors.[28] Our interests and beliefs are embedded in the vocabulary we use, the questions we ask, perhaps our whole search history. YouTube, Facebook, and other social media present content based on what we have wanted to see in the past, and what other people who are like us across some array of dimensions have wanted to see.

In this way, platform companies have become knowledge intermediaries, like newspapers or school curriculum boards, while insulating themselves from traditional accountability. Their algorithms and (perhaps just as important) sharing and search tools help foster categories that can become self-reinforcing private universes of discourse, producing echo chambers in which other voices are silenced, or epistemic bubbles that guide users to apparent authorities who actively look to discredit other sources of information.[29] However, the invisibility of hier-

archy allows these knowledge intermediaries to justify themselves on *laissez-faire* principles, not telling the public what to trust, even while they quietly sink deeper into the Augean mire of moderating offensive, false, or untrue content.

Our universe of accessible knowledge is shaped by categorization processes that are invisible and incomprehensible to ordinary users, according to principles that have little regard for whether it is well sourced. The outcome is that the way that people "take [their] bearings in the world" is slowly changing.[30] Visible feedback loops between the people being categorized, the knowledge they have access to, and the processes through which the categories are generated are replaced by invisible loops mediated through algorithms that maximize on commercial imperatives, sometimes creating incompatible and self-sustaining islands of shared ("post-truth") beliefs among micropublics who have been categorized in particular ways, and who may themselves act to reinforce the categories. A new terrain of political struggle has arisen, involving the exploitation of information systems and algorithmic dynamics for partisan advantage.

This is a different set of moral pathologies than those suggested by social psychologist Shoshana Zuboff, who emphasizes platform companies' manipulation of people's wants and beliefs, which might or might not succeed.[31] The more corrosive threat may be that people have been convinced that the high-tech modernist system of knowledge generation is an open buffet where "anything goes," and that keeping it that way is essential to their own freedom. Anyone can offer content, anyone can be their own expert, and it is up to the algorithm to sort it out. Further, the new existential condition of transparency has provided everyone with potent tools to expose or doubt others, only moderated by their own vulnerability to be exposed in turn – an inherently agonistic situation.

At the end of the day, the relationship between high modernism and high-tech modernism is a struggle between two elites: a new elite of coders, who claim to mediate the wisdom of crowds, and an older elite who based their claims to legitimacy on specialized professional, scientific, or bureaucratic knowledge.[32] Both elites draw on rhetorical resources to justify their positions; neither is disinterested.

The robust offense and disbelief that many people feel about algorithmic judgments suggests that the old high modernist moral political economy, faults and all, is not quite dead. The new moral political economy that will replace it has not yet matured, but is being bred from within. Articulated by technologists and their financial backers, it feeds in a kind of matriphagy on the enfeebled body (and the critique) of its progenitor. Just as high modernist bureaucracies did before, high-tech modernist tools and their designers categorize and order things, people, and situations. But they do so in distinctive ways. By embedding surveillance into everything, they have made us stop worrying about it, and perhaps even come

to love it.[33] By producing incomprehensible bespoke categorizations, they have made it harder for people to identify their common fate. By relying on opaque and automated feedback loops, they have reshaped the possible pathways to political reaction and resistance. By increasing the efficiency of online coordination, they have made mobilization more emotional, ad hoc, and collectively unstable. And by insisting on market fairness and the wisdom of crowds as organizing social concepts, they have fundamentally transformed our moral intuitions about authority, truth, objectivity, and deservingness.

ABOUT THE AUTHORS

Henry Farrell is the SNF Agora Professor of International Affairs at Johns Hopkins School of Advanced International Studies. He is the author, with Abraham Newman, of *Underground Empire: How America Weaponized the World Economy* (forthcoming 2023) and *Of Privacy and Power: The Transatlantic Fight over Freedom and Security* (2019).

Marion Fourcade is Professor of Sociology at the University of California, Berkeley. She is the author of *Economists and Societies: Discipline and Profession in the United States, Britain, and France, 1890s to 1990s* (2010) and editor of *Pandemic Exposures: Economy and Society in the Time of Coronavirus* (with Didier Fassin, 2022).

ENDNOTES

  [1] Mary Douglas, *How Institutions Think* (Syracuse, N.Y.: Syracuse University Press, 1986), 91.

  [2] Langdon Winner, "Do Artifacts Have Politics?" *Dædalus* 109 (1) (Winter 1980): 121–136.

  [3] Virginia Eubanks, "The Mythography of the 'New' Frontier," MIT Communications Forum, 1999, https://web.mit.edu/comm-forum/legacy/papers/eubanks.html.

  [4] James Scott, *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed* (New Haven, Conn.: Yale University Press, 1998).

5 Robyn Caplan and danah boyd, "Isomorphism through Algorithms : Institutional Dependencies in the Case of Facebook," *Big Data & Society* 5 (1) (2018): 1–12.

6 William Cronon, *Nature's Metropolis : Chicago and the Great West* (New York : W. W. Norton, 1991).

7 Marion Fourcade and Kieran Healy, "Seeing Like a Market," *Socio-Economic Review* 15 (1) (2017): 9–29.

8 Luke Stark, "The Emotive Politics of Digital Mood Tracking," *New Media and Society* 22 (11) (2020): 2039–2057.

9 James Scott, *Weapons of the Weak : Everyday Forms of Peasant Resistance* (New Haven, Conn. : Yale University Press, 1985).

10 Ian Hacking, *The Social Construction of What?* (Cambridge, Mass. : Harvard University Press, 1999), 103–104.

11 Geoffrey Bowker and Susan Leigh Star, *Sorting Things Out : Classification and Its Consequences* (Cambridge, Mass. : The MIT Press, 1999).

12 Georges-Eugène Haussmann was the prefect responsible for the renewal and reimagining of Paris in Napoleonic France.

13 Cathy O'Neil, *Weapons of Math Destruction : How Big Data Increases Inequality and Threatens Democracy* (New York : Crown, 2016); and Jenna Burrell, "How the Machine 'Thinks' : Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* 3 (1) (2016): 1–12.

14 Martha Poon, "From New Deal Institutions to Capital Markets : Commercial Consumer Risk Scores and the Making of Subprime Mortgage Finance," *Accounting, Organizations and Society* 34 (5) (2009): 654–674.

15 Greta Krippner, "Democracy of Credit : Ownership and the Politics of Credit Access in Late Twentieth-Century America," *American Journal of Sociology* 123 (1) (2017): 1–47.

16 Kai-Fu Lee, *AI Superpowers : China, Silicon Valley and the New World Order* (New York : Harper Business, 2018).

17 Solon Barocas and Andrew D. Selbst, "Big Data's Disparate Impact," *California Law Review* 104 (3) (2016): 671–732; Ruha Benjamin, *Race After Technology : Abolitionist Tools for the New Jim Code* (Cambridge, Mass. : Polity, 2019); and Safiya Umoja Noble, *Algorithms of Oppression : How Search Engines Reinforce Racism* (New York : NYU Press, 2018).

18 European Commission, *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts* (Brussels : European Commission, 2021).

19 Angèle Christin, *Metrics at Work : Journalism and the Contested Meaning of Algorithms* (Princeton, N.J. : Princeton University Press, 2020); and Sarah Brayne, *Predict and Surveil : Data, Discretion, and the Future of Policing* (New York : Oxford University Press, 2020).

20 Barbara Kiviat, "The Moral Limits of Predictive Practices : The Case of Credit-Based Insurance Scores," *American Sociological Review* 84 (6) (2019): 1134–1158.

21 Marion Fourcade, "Ordinal Citizenship," *The British Journal of Sociology* 72 (2) (2021): 154–173.

22 Peter Thiel, "Competition Is for Losers," *The Wall Street Journal*, September 12, 2014.

23 Scott, *Seeing Like a State*.

24 Gil Eyal, *The Crisis of Expertise* (Cambridge, Mass.: Polity, 2019).

25 John Perry Barlow, "A Declaration of the Independence of Cyberspace," Electronic Frontier Foundation, February 8, 1996, https://www.eff.org/cyberspace-independence; Friedrich von Hayek, "The Uses of Knowledge in Society," *American Economic Review* 35 (4) (1947): 519–530; Friedrich von Hayek, "Competition as a Discovery Procedure," *The Quarterly Journal of Austrian Economics* 5 (3) (2002): 9–23; and Evgeny Morozov, "Digital Socialism? The Socialist Calculation Debate in the Age of Big Data," *New Left Review* 116/117 (2019), https://newleftreview.org/issues/ii116/articles/evgeny-morozov -digital-socialism.

26 Wendy Brown, *Undoing the Demos: Neoliberalism's Stealth Revolution* (Princeton, N.J.: Princeton University Press, 2015), 157.

27 Lina M. Khan, "Amazon's Antitrust Paradox," *Yale Law Journal* 126 (3) (2016–2017): 710–805.

28 Francesca Tripodi, *Searching for Alternative Facts. Analyzing Scriptural Inference in Conservative News Practices* (New York: Data & Society, 2018), https://datasociety.net/library/ searching-for-alternative-facts.

29 C. Thi Nguyen, "Echo Chambers and Epistemic Bubbles," *Episteme* 17 (2) (2020): 141–161.

30 Hannah Arendt, "Truth and Politics," in *The Portable Hannah Arendt*, ed. Peter Baehr (London: Penguin Classics, 2000), 568.

31 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: Public Affairs, 2019); and Tim Hwang, *Subprime Attention Crisis* (New York: FSG Originals, 2020).

32 William Davies, "Elite Power Under Advanced Neoliberalism," *Theory, Culture and Society* 34 (5–6) (2017): 227–250; and Jenna Burrell and Marion Fourcade, "The Society of Algorithms," *Annual Review of Sociology* 47 (2021): 213–237.

33 Nitsan Chorev, "The Virus and the Vessel, or: How We Learned to Stop Worrying and Love Surveillance," *Socio-Economic Review* 19 (4) (2021): 1497–1513.