2019

Privacy's Political Economy and the State of Machine Learning

Mariano-Florentino Cuéllar

Aziz Z. Huq

Privacy's Political Economy and the State of Machine Learning: An Essay in Honor of Stephen J. Schulhofer

(forthcoming 2020 in *NYU Annual Survey of American Law*)

Mariano-Florentino Cuéllar* and Aziz Z. Huq**

Abstract

Our aim in this essay is to consider how policymakers make decisions about government surveillance in what we might call the machine learning state — a nation-state equipped with sufficient bureaucratic and technological capacity to rely extensively on machine learning techniques for surveillance, law enforcement, and national security. We focus particularly on the question of how the state's political economy influences its decision to adopt privacy-relevant machine learning technologies. Since machine learning tools can also be deployed in many ways that are not pertinent to privacy, our focus therefore is on a specific subset of state uses of such technology — to engage in surveillance of the public and its activities. In order to lay the groundwork for nuanced engagement with the legal and policy trade-offs in this domain, we aim here map the main technological and institutional forces shaping a state's deployment of new machine learning capabilities that can affect privacy, and then to explore, more tentatively, their likely effects on technological uptake.

We propose that state adoption of machine learning instruments for surveillance occurs in a variation on what Robert Putnam famously characterized as a "two-level game." The state is operating simultaneously in a domestic political environment populated by institutions mediating conflicts involving civil society and firms competing to expand and monetize machine learning capacities, and also in an international environment in which it is competing with other sovereign nations that are cultivating and deploying similar capacities for geostrategic ends. How and to what end machine learning instruments are deployed depends on the strategic choices that the national government makes in these two overlapping yet distinct contexts. We emphasize that it would be a mistake to analyze these choices purely in terms of responses to domestic considerations, because states may be willing to shoulder the costs of domestic backlash so they can further geostrategic goals. We also elucidate the way the net vector of domestic and international pressures is likely to affect legal and policy interventions in this space. Based on this exercise, we identify concerns and trade-offs relevant to the possible reforms; and explore both the limitations of existing Fourth Amendment doctrine, and the potential of state and federal legislative or regulatory alternatives as reform instruments.

Introduction

In September 2002, Stephen Schulhofer published a short book entitled "The Enemy Within: Intelligence, Law Enforcement and Civil Liberties after September 11." Those expecting a breathless "libertarian panic"[1] from a leading liberal scholar of constitutional criminal procedure would have been disappointed. Instead "The Enemy Within" reflects Schulhofer's characteristically scrupulous care and lawyerly skill. Taking meticulous care to sidestep harsh overreaction, he sets out to map the

* Justice, California Supreme Court, Herman Phleger Visiting Professor of Law, Stanford Law School, and affiliated scholar, Freeman Spogli Institute for International Studies at Stanford University.

** Frank and Bernice J. Greenberg Professor of Law and Mark C. Mamolen Teaching Scholar, University of Chicago Law School. Thanks to Kristen Eichensehr for useful feedback, and the Frank J. Cicero Fund for support.

[1] STEPHEN SCHULHOFER, THE ENEMY WITHIN (2002)("Enemy Within"). A skeptical neologism coined in Adrian Vermeule, *Libertarian Panics*, 36 RUTGERS L.J. 871, 873 (2005). The empirical premises of the concept are, however, fragile. Aziz Z. Huq, *Structural Constitutionalism as Counterterrorism*, 100 CALIF. L. REV. 887, 934-43 (2012).

complex practical, statutory, and doctrinal terrain of domestic electronic surveillance—how information-gathering techniques justified on national security grounds were not confined to the national security sphere, for example,[2] and how the state's information-gathering activities are inevitably mediated by organizational practices and institutional realities.[3] By subjecting this terrain to careful scholarly attention, he demonstrated how to glean from its features some degree of normative guidance. One could pick any number of Schulhofer's works to make similar observations. But "The Enemy Within" provides an especially salient launching point for this essay. It not only offers important historical groundwork for our project, it also evinces a more general methodological orientation relevant to scholars working on fractious national security, technology, and privacy related questions.

In the roughly two eventful decades since publication of "The Enemy Within," the ensuing political and technological changes heighten the importance of new inquiries into the circumstances under which people can make privacy claims of diverse sorts against the state.[4] One set of changes concerns the emergence and deployment at scale of new computational tools—often described using terms like 'machine learning' or 'artificial intelligence—to extract correlations or predictive inferences from large data sets. Advocates and critics of these technologies alike make bold claims about the extension of epistemic capabilities flowing from new computational instruments.[5] Already, machine learning tools underline preductions criminal violence and play some role in allocating responsive state coercion.[6] Increasingly, they also play a role in the allocation of military force.[7] Deployment of new computational technologies in the criminal justice context has provoked concerns about their "threat[]" to privacy values.[8] To date, however, we still lack an account of the social, economic, and political forces that shape adoption of these technologies *by (or on behalf of) the state*.[9] Without a coherent account of this political economy, however, we are ill-equipped to evaluate either the justifications or the likely trajectory of such technological change in the forms of state power. We'd also lack a crucial tool to identify meaningful efforts, or the ultimate consequences for privacy in relation to the state.

Our aim in this essay is to offer a preliminary sketch of the basic political economy of privacy in what we might call *the machine learning state*. This is a nation state with sufficient bureaucratic and

---

[2] *Enemy Within*, supra note __, at 1-7, 29-48.

[3] *Id.* at 55-64.

[4] We bracket here the question whether privacy's principal adversary is no longer the state, but instead the coterie of companies that harvest and analyze personal data. We aim to take up that question in other work.

[5] *Compare* PEDRO DOMINGOS, THE MASTER ALGORITHM: HOW THE QUEST FOR THE ULTIMATE LEARNING MACHINE WILL REMAKE OUR WORLD 1 (2015) (offering an optimistic take on machine learning's impact), *with* CATHY O'NEILL, WEAPONS OF MATH DESTRUCTION: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY 203–06 (2016) (decrying the regressive tendencies of big-data technologies generally).

[6] Aziz Z. Huq, *Racial Equity in Algorithmic Criminal Justice*, 109 DUKE L.J. 1043, 1068-76 (2019) (documenting adoption of new risk-assessment algorithms).

[7] Ashley S. Deeks, *Predicting Enemies*, 104 VA. L. REV. 1529 (2018).

[8] Kiel Brennan-Marquez, *"Plausible Cause": Explanatory Standards in the Age of Powerful Machines*, 70 VAND. L. REV. 1249, 1255-57 (2017); Emily Berman, *A Government of Laws and Not of Machines*, 98 B.U. L. REV. 1277, 1339 (2018) (exploring privacy concerns).

[9] In other work, we have criticized the leading extant accounts as insufficiently attentive to the effects of state adoption of such tools, *See* Mariano-Florentino Cuéllar and Aziz Z. Huq, *Economies of Surveillance*, 133 HARV. L. REV. 1280 (2020).

technological capacity to rely extensively on machine learning techniques for surveillance, enforcement, and security. We focus particularly in on the question of how the machine learning state's political economy influences its decisions to adopt privacy-relevant technologies. Since machine learning tools can be deployed in many ways not pertinent to privacy, our focus is on a narrow subset of such deployments. We make no claim to canvas the waterfront of machine learning as an instrument of state policy. Though we begin mapping some of the normative trade-offs that arise for such states and their citizens, our attention here is largely taken up by descriptive questions, not prescriptive ones. To lay the groundwork for more nuanced engagement with the legal and policy trade-offs in this domain, we map what we perceive to be the main technological and institutional forces that shape state's adoption of new machine learning instruments.[10] Even if we acknowledge that people, organizations, and states are driven by a variety of motivations, the more general pressures and incentives we incorporate into our political-economy framework are pertinent in that they help us map the possibility conditions of privacy vis-à-vis the state in the age of machine learning.

Our central contribution here is to suggest that state adoption of machine learning instruments for surveillance occurs in a version of what Robert Putnam famously characterized as a "two-level game."[11] That is, the state is operating simultaneously in a domestic political environment dominated by firms competing to expand and monetize machine learning capacities, and simultaneously in an international environment in which it is competing with other sovereign nations that are cultivating and deploying the same technological capacities for geostrategic ends. How and to what end machine learning instruments are deployed turns on the strategic choices that the national government makes in these two overlapping yet distinct contexts.[12] It would thus be a mistake to analyze such deployment in terms of responses to purely domestic legal, policy, or political considerations.

Although our focus is primarily descriptive, we also offer a preliminary sense of how to think about the possibility conditions for privacy in the context of this two-level game. We also recognize the limitations of existing Fourth Amendment doctrine as a means of addressing the situation we describe and the potential for legislative and regulatory alternatives to fill the gap, though we offer no fully-developed prescriptions at this stage. Rather, we hope to elucidate the pressures likely to face legal and policy interventions in this space. In the process, we aim to map some of the concerns and trade-offs that reasonable interventions would need to address. In developing this account, we make no claim that machine learning is not the only technology whose deployments are shaped by

[10] We focus on the federal government. State and municipal governments adopt predictive computational tools for different purposes, under different fiscal constraints, and under different political conditions. For a recent account of local resistance to national control of local security functions, see Trevor G. Gardner, *Immigrant Sanctuary as the 'Old Normal': A Brief History of Police Federalism*, 119 COLUM. L. REV. 1 (2019).

[11] Robert D. Putnam, *Diplomacy and domestic politics: the logic of two-level game*s, 42 INT'L ORG. 427 (1988). Our use of Putnam analytic framework, we stress, is limited. We are not focused, for example on his account of negotiation dynamics. Rather, it is the possibility of dynamic interactions between the domestic and the international that we find most useful.

[12] Our analytic frame is consistent with the "new interdependence approach" developed by Henry Farrell and Abraham Newman, in which "institutions [act] as opportunity structures that facilitate cross-national coordination between collective actors" and (particularly resonant with our account) "political contestation … takes place in multiple and overlapping venues." HENRY FARRELL AND ABRAHAM L. NEWMAN, OF PRIVACY AND POWER: THE TRANSATLANTIC STRUGGLE OVER FREEDOM AND SECURITY 29 (2019).

interleaved domestic and international dynamics game.[13] Quite the contrary. But we think there are distinctive ways in which similar dynamic influences how machine learning specifically is adopted.

In Part I, we define "machine learning" and "artificial intelligence," and explain how such technologies tend to have implications (negative and positive) for various forms of privacy. Part II develops the claim that adoption of such technologies occurs against the context of a two-level domestic and international game. It further considers how that complex context might shape the privacy impacts of machine learning, thereby offering the essential context for closely-related related doctrinal questions and institutional design problems. A conclusion briefly considers how those impacts, to the extent that they are perceived as undesirable, might be mitigated.

I. Machine Learning's Impact on Privacy

As terms like "machine learning" and "privacy" are far from self-explanatory, it is is useful to begin by offering definitions of these two central concepts. We can then offer a preliminarily sketch of the technologies' implications for privacy by first noting how privacy intrusions might occur, and then explaining how emerging technology can enable or constrain such intrusions.

Exactly what is enabled by any technological change is contingent on the organizational and legal context shaping society's use and understanding of such change. The technical possibility that computational tools can impinge on privacy does not mean that such harms will inevitably arise. To the contrary, one of our key assumptions is that the manner in which new technologies are adopted depends on social as well as technical conditions. A "successful technological innovation occurs only when all the elements of the system, the social as well as the technological, have been modified so that they work together effectively."[14] The immanent potentialities of a new technology remain unexpressed in the absence of social, institutional, and economic circumstances (or "affordances")[15] in which they become relevant. Our account of the 'raw' technology of machine learning, therefore, must be understood as necessarily incomplete, its implications latent and unexpressed until revealed by social and institutional context.

A. The Domain of Machine Learning

We are concerned in this essay with a group of computational tools called "machine learning." Machine learning, in its most general terms, is a technique for using computing platforms to solve a "learning problem . . . [for] improving some measure of performance when executing some task through some type of training experience."[16] Although they are often called "artificial intelligence" tools, we generally avoid that term here because—relative to the phenomena we describe here—artificial intelligence may be both somewhat over-broad and under-inclusive relative to the primary

[13] *See id.* at 69-160 (developing detailed case studies of interjurisdictional conflict about financial and airline passenger data).

[14] Bryan Pfaffenberger, *Social anthropology of technology*, 21 ANN. REV. ANTHROPOLOGY 491, 498 (1992).

[15] Ryan Calo, *Can Americans Resist Surveillance?*, 83 U. CHI. L. REV. 23 (2016) (using the term "affordance" to describe not only the enabling effects of artifacts but also those of cultural practices and norms).

[16] M.I. Jordan & T.M. Mitchell, *Machine Learning: Trends, Perspectives, and Prospects*, 349 SCI. 255, 255 (2015).

focus of our analysis.[17] Whether or not the computational tools we have in mind here are capable of self-modification through "learning" in the way the "artificial intelligence" label might suggest, they are invaluable in analyzing vast quantities of data that would be enormously cumbersome for humans to sift, and are capable of identifying subtle patterns that could elude even perceptive human observers or analysts using only conventional tools of statistical inference.

In most respects, the basic intuition animating the modality of machine learning is not new. An elementary form of the underlying computational model, called the perceptron, was developed to facilitate supervised learning of binary classifiers; it's been well understood since the late 1950s.[18] An important technical breakthrough, however, occurred in 1985, when the computer scientists Geoffrey Hinton developed a tool called 'backpropagation.' This enabled a spectacular and rapid adoption of a kind of machine learning called "neutral networks." Nevertheless, it would take another 26 years before sufficient computing power was generally available to make Hinton's method a plausible one for commercial use.[19] A brief explanation of the basic technology helps us understand its range of possible state uses.

Most machine learning algorithms ordinarily work by sorting a class of examples (e.g., images or individuals) into a set of categories.[20] For instance, a machine learning tool for visual recognition might sort images into the classes of "face" and "not face." A bail algorithm might class suspects into the classes of "very dangerous," "dangerous," and "not dangerous." The classification is possible because the algorithm has already encountered a set of training data—i.e., a data set in which the individual items have already been classified. By examining relationshps between individuals in the training data and an outcome of interest, the algorithm can "learn[] rules [for classification] from data," rules that can then be applied to new, previously unknown data sets.[21] The ensuing classifications of new data, however, are typically correlational and not causal in nature. As a result, a machine learning algorithm's performance is usually gauged in terms of how well it captures the strength of the relation of x to y, and not by its ability to discern an actual causal relationship of x and y.[22]

Machine learning can be used to make predictions of outcomes in the case of supervisory data, or to identify clusters or associations (in the case of recommendation systems such as those employed

---

[17] *Cf.* STUART RUSSELL & PETER NORVIG, ARTIFICIAL INTELLIGENCE: A MODERN APPROACH 2-14 (3d ed. 2013) (offering a series of alternative definitions of AI that encompass concepts and processes different from what "machine learning" covers, include thinking and acting humanly as well as rationally). The international relations literature uses the term 'artificial intelligence,' so we find we cannot completely avoid it.

[18] Frank Rosenblatt, *The perceptron: a probabilistic model for information storage and organization in the brain*, 65 PSYCH. REV. 386 (1958).

[19] James Somers, *Is AI Riding a One-Trick Pony*?, 120 MIT TECH. REV. 29, 31 (2017).

[20] Comm. on the Analysis of Massive Data et al., *Frontiers in Massive Data Analysis* 104 (2013), http://www.nap.edu/catalog.php?record_id=18374; *accord* PETER FLACH, MACHINE LEARNING: THE ART AND SCIENCE OF ALGORITHMS THAT MAKE SENSE OF DATA 14 (2012).

[21] Ziad Obermeyer & Ezekiel J. Emanuel, *Predicting the future—big data, machine learning, and clinical medicine*, 13 NEW ENGLAND J. MED. 1216, 1217 (2016); PEDRO DOMINGUES, THE MASTER ALGORITHM: HOW THE QUEST FOR THE ULTIMATE LEARNING MACHINE WILL REMAKE OUR WORLD 6-7, 23 (2015).

[22] Jordan & Mitchell, *supra* note 16, at 255-57 (noting that performance can be defined in terms of accuracy, with false positive and false negative rates being assigned a variety of weights).

by Netflix and Amazon).[23] Computationally, machine learning tools can be implemented through a wide range of strategies. These include associational learning,[24] 'neural networks,'[25] and the "random forests" approach.[26] Random forests, and the wider category of decision-tree models in which they fall, are particularly useful for nominal or ordinal data; in contrast, neural networks work well with numerical data.[27] We will ignore the differences between these approaches for present purposes.

One taxonomy of machine learning's practical applications employs four categories of functionalities: (i) the identification of clusters or associations, within a population under analysis; (ii) the identification of outliers within a population; (iii) the development of associational rules; and (iv) prediction problems of classification and regression applied to out-of-sample data.[28] Examples abound, and we offer only a small handful here to illustrate the technology's possibilities. A first comes a recent study of the allocation of hip replacement surgery among otherwise eligible patients.[29] The study used machine learning tools to identify which patients would live long enough to benefit from the surgery.[30] A second use involves algorithms designed to predict the spatial occurrence of future crime patterns, and hence to help determine future police deployments.[31] A third is the use of machine learning instruments to scan large volumes of video footage in search of specific faces. In the United Kingdom, for instance, facial recognition tools have been used since 2016 to identify suspects from surveillance and make arrests.[32] Law enforcement in the United States is presently adopting a similar tool, to some controversy given the absence of oversight over that its roll-out and uncertainty and the instrument's quality.[33] The list of public and private uses could be extended for some time without running out of extraordinary and novel uses of the technology.

B. Machine Learning against Privacy

---

[23] Judea Pearl, *Theoretical impediments to machine learning with seven sparks from the causal revolution*, at 1-2 (2018) (arguing that inability of machine learning to analyze counterfactuals to infer causation is a major impediment).

[24] Trevor Hastie, Robert Tibshirani & Jerome Friedman, *Unsupervised Learning*, in THE ELEMENTS OF STATISTICAL LEARNING 485 (2009).

[25] The following draws on the lucid accounts in ETHEM ALPAYDIN, INTRODUCTION TO MACHINE LEARNING 88-103 (3d ed. 2014), and Yoshua Bengio, *Machines Who Learn*, SCIENTIFIC AM., June 2016, at 46; *see also* SEAN GARISH, HOW SMART MACHINES THINK 109-23 (2019) (providing a lucid explanation of neural networks in action).

[26] Leo Breiman, *Random forests*, 45 MACHINE LEARNING 5, 5 (2001).

[27] *Id.* at 136.

[28] JOHN D. KELLEHER & BRENDAN TIERNEY, DATA SCIENCE 151-80 (2018) (providing examples of these different tasks).

[29] Jon Kleinberg et al., *Prediction policy problems*, 105 AM. ECON. REV. 491, 493-94 (2015).

[30] *Id.* at 493.

[31] *See* Laura Nahmias & Miranda Neubauer, *NYPD Testing Crime-Forecast Software*, POLITICO (July 8, 2015, 5:52 AM), https://www.politico.com/states/new-york/city-hall/story/2015/07/nypd-testing-crime-forecast-software-090820 [https://perma.cc/3G49-UP9B].

[32] Cara McGoogan, *British Police Arrest Suspect Spotted with Facial Recognition Technology*, TELEGRAPH (June 7, 2017, 4:31 PM), http://www.telegraph.co.uk/technology/2017/06/07/british-police-arrest-suspect-spotted-facial-recognition-technology [https://perma.cc/Q5H7-W4QF]; *see generally* Sarah Brayne, *Big Data Surveillance: The Case of Policing*, 82 Am. Soc. Rev. 977, 980 (2017) (documenting the adoption of such technologies by the Los Angeles Police Department).

[33] Kashmir Hill, *The Secretive Company that might end Privacy as we know it*, N.Y. TIMES (Jan. 18, 2020), https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html.

To further understand the implications of machine learning for privacy, we must acknowledge the often-contentious debates associated with even defining, let alone weighing the trade-offs associated with protecting, privacy. Well-before privacy became entangled in late twentieth century doctrinal disputes about federal constitutional rights, what Samuel Warren and Louis Brandeis called "the right to be let alone" had already begun catalyzing spirited disagreement about matters such as the scope of civic life, and the nature of dignity.[34] Privacy sparked controversy essentially since it emerged as a subject of discussion in American political culture in the late nineteenth century.[35] The emergence of Fourth Amendment jurisprudence organized around the idea of privacy, which happened only in the second half of the twentieth century, did not abate that controversy. To the contrary, the jurisprudence still encompasses a jostling bundle of analytically distinct models of privacy,[36] while scholars still argue about appropriate underlying theoretical model of Fourth Amendment protection in terms that either reinterpret or reject a privacy touchstone.[37]

Without trying to settle these seemingly intractable debates, we identify forms of privacy that might be implicated by the adoption of machine learning tools, and where possible supply examples. We map these forms of privacy using examples from the Fourth Amendment case-law, although we do not intend for our analysis to begin and end with that body of constitutional doctrine. Rather, the cases are is simply helpful as a source of illustration.

The first, and perhaps most obvious, form of privacy impacted by machine learning tools is 'informational' privacy, i.e., a person's interest in preventing the disclosure of information that she wishes to keep secret. This form of privacy might be conceptualized as a kind of intellectual property interest: a right to control the possibility of transactions over or dissemination of a given piece of information. For instance, in the recent decision of in *Carpenter v. United States*, the U.S. Supreme Court described government acquisition of cell-site locational data from a suspect's telecommunications provider as "a new phenomenon: the ability to chronicle a person's past movements through the record of his cell phone signals."[38]

A second potential form of privacy focuses on its dignitary aspect. This might be understood to encompass instances in which a physical space, such as a home, is viewed as a domain of special normative concern such that intrusions on that space are perceived to be undesirable, even objectionable, without regard to whether they yield the disclosure of any new information. Hence, the

---

[34] Samuel Warren and Louis D. Brandeis, *The Right to Privacy*, 4 HARV. L. REV. 193, 504-08 (1890).

[35] SARAH E. IGO, THE KNOWN CITIZEN: A HISTORY OF PRIVACY IN MODERN AMERICA 2-3 (2018).

[36] *See* Orin S. Kerr, *Four Models of Fourth Amendment Protection*, 60 STAN. L. REV. 503, 504-08 (2007) (explaining Fourth Amendment doctrine in terms of four competing paradigms).

[37] For a reinterpretation of privacy, see, e.g., Matthew B. Kugler & Lior Jacob Strahilevitz, *Actual Expectations of Privacy, Fourth Amendment Doctrine, and the Mosaic Theory*, 2015 SUP. CT. REV. 205, 211 (including "perceived intrusiveness" of a search as relevant to reasonable expectations of privacy per Fourth Amendment doctrine); and Richard M. Re, *Fourth Amendment Fairness*, 116 MICH. L. REV. 1409, 1413 (2018) (contending "a search or seizure is unreasonable when any principle that permitted it would be one that a Fourth Amendment rights holder could reasonably reject"). For a rejection of a privacy touchstone, see William Baude & James Y. Stern, *The Positive Law Model of the Fourth Amendment*, 129 HARV. L. REV. 1821, 1877 (2016) (reasoning that, under the positive law model, a court may decide to apply the waiver of positive law rights to Fourth Amendment protections for threshold search and seizure questions).

[38] 138 S. Ct. 2206, 2216 (2018).

U.S. Supreme Court has repeatedly described the home as a location deserving of special solicitude under the Fourth Amendment.[39] The home is protected doctrinally without regard to whether the information secured through an intrusion therein yields information.[40]

Yet a third dimension of privacy relates to political power, and the nature of the relationship between the state and those subject to its authorities. Privacy, that is, is a political concept as well as a legal or ethical concept insofar as it speaks to the nature of the relationship between (say) the state and the individual subjects (citizen and noncitizen) that it potentially regulates. Privacy is a method of calibrating the distance between state and subject so as to maintain a certain "equilibrium" defined in terms of the power the state potentially exercises over the subject.[41] This political notion is consistent with the origins of the Fourth Amendment (although perhaps not with the manner in which it has now been implemented in the quotidian criminal law context). At its inception, that Amendment "was about maintaining space for individuals to compete for offices created by the separation of powers system--individuals who might play vital roles in resisting incipient despotism."[42]

In the hands of a state functional enough to staff and deploy complex organizations, machine learning in state hands has the potential to influence each of these three forms of privacy. Consider first its possible impact on informational privacy. Machine learning tools can be applied to large pools of publicly available data in order to acquire information that otherwise would not be available. For example, the metadata from telephone communications can be analyzed without machine learning to "reveal[] what and who we're interested in and what's important to us no matter how private," including illnesses (both ours and those of people close to us); intimate decisions (such as decisions to either have children or to abort a pregnancy); and decisions to acquire goods, such as a weapon.[43] It has long been the case that a diligent investigator could infer some things from aggregated transactional records, say from a bank or telephone company. Machine learning, however, makes the analysis of large pools of such transactional data much more revealing. Hence, one study has used such tools to extract both age and gender information solely from metadata about telephone usage.[44] This change in the magnitude of inference might translate into a significant change in the sheer power of the state.

Further, machine learning can be used to expand the range of data that is epistemically fruitful, thereby allowing for inferences in ways that compromise both informational privacy and dignity values. One well-known example involves the de-anonymization of large putatively non-individualized

---

[39] Wilson v. Layne, 526 U.S. 603, 612 (1999) (describing protection of the home as "core" to the Fourth Amendment); Stephanie M. Stern, *The Inviolate Home: Housing Exceptionalism in the Fourth Amendment*, 95 CORNELL L. REV. 905, 912-13 (2010).

[40] An example of this is the rule requiring officers to knock and announce their presence when executing a warrant. Wilson v. Arkansas, 514 U.S. 927 (1995). The so-called 'knock and announce' rule does not prevent the disclosure of information. It does protect a dignitary interest in the home.

[41] Orin S. Kerr, *An Equilibrium-Adjustment Theory of the Fourth Amendment*, 125 HARV. L. REV. 476, 480 (2011).

[42] Aziz Z. Huq, *How the Fourth Amendment and the Separation of Powers Rise (and Fall) Together*, 83 U. CHI. L. REV. 139, 146–47 (2016)

[43] BRUCE SCHNEIER, DATA AND GOLIATH: THE HIDDEN BATTLES TO COLLECT YOUR DATA AND CONTROL YOUR WORLD 24-25 (2015).

[44] B. Felbo et al., *Using deep learning to predict demographics from mobile phone metadata* (2016), https://openreview.net/forum?id=91EEnoZX0HkRlNvXUKLA.

datasets.[45] Another is the use of data generated by internet usage to profile and predict behavior—a measure that is already standard among private advertisers[46]—and underlying psychological states. A 2017 study, for instance, showed how a trained algorithm could use Instagram feeds to predict markers of clinical depression better than a human diagnosis.[47] Computer vision can also reveal information that otherwise could not be secured—such as the identity of a person being extracted from visual data solely depicted that person's gait.[48] The IC Realtime Company now offers an application called "Ella," which can recognize and execute natural language queries for CCTV footage.[49] The outer perimeter of computational infererence remains analytically and ethically murky. In 2018, a pair of Stanford researchers published research controversially suggesting that sexual preferences could be accurately inferred from facial images.[50] Their findings, however, have since been challenged on technical grounds, although it remains unclear whether inference of traits such as sexual preference from facial images will be feasible in the near term.[51] Indeed, under the right circumstances (say, in the hands of a very conservative or theocratic state), even a weakly predictive instrument for drawing sexuality-related preferences from image data might pose significant and troubling normative implications. And having the state make imprecise predictions about sexuality, we hasten to add, can impinge on individual dignity in troubling ways even in a well-functioning democracy. It is also important to note that at least some uses of facial recognition technology work simply as substitutes for presently available technologies in ways that have no discernable privacy impact. The use of facial recognition to manage secure entry of government employees into their workplace, for instance, in lieu of other forms of identification, for instance, does not provide a reason for new privacy related concern—unless, of course one has a functional theory addressing how deployment of such technology in the building access context might facilitate public habituation or doctrinal justifications for more

[45] The pathbreaking example involved a de-anonymization of the Netflix database used in its algorithm design contest and credit card data. Arvind Narayanan & Vitaly Shmatikov, *Robust De-anonymization of Large Sparse Datasets*, *in* PROCEEDINGS OF THE 2008 IEEE SYMPOSIUM ON SECURITY AND PRIVACY 111-125 (2008) ("[V]ery little auxiliary information is needed [to] deanonymize an average subscriber record from the Netflix Prize dataset."); Yves-Alexandre de Montjoye et al., *Unique in the Shopping Mall: On the Reidentifiability of Credit Card Metadata*, 347 Sci. 536 (2015) ("We study 3 months of credit card records for 1.1 million people and show that four spatiotemporal points are enough to uniquely reidentify 90% of individuals."). Where other databases are available to be cross-references, it may be even easier to de-anonymize data. Latanya Sweeney, *k-Anonymity: A Model For Protecting Privacy*, 10 INT'L J. ON UNCERTAINTY, FUZZINESS & KNOWLEDGE-BASED SYS. 557, 557 (2002).

[46] Sonia K. Katyal, *Private Accountability in the Age of Artificial Intelligence*, 66 UCLA L. REV. 54, 91 (2019) ("Since websites often rely on predictive algorithms to analyze people's online activities--web surfing, online purchases, social media activities, public records, store loyalty programs, and the like--they can create profiles based on user behavior, and predict a host of identity characteristics that marketers can then use to decide the listings that a user sees online.").

[47] Andrew G. Reece, & Christopher M. Danforth, *Instagram photos reveal predictive markers of depression*, 6 EPJ DATA SC. 15, 15-16 (2017).

[48] SCHNEIER, *supra* note 43, at 34.

[49] James Vincent, *Artificial Intelligence is Going to Supercharge Surveillance*, THE VERGE (Jan. 23, 2018, 10:54 AM), https://www.theverge.com/2018/1/23/16907238/artificial-intelligence-surveillance-cameras-security [https://perma.cc/U4SN-VFL2].

[50] Yilun Wang and Michal Kosinski, *Deep neural networks are more accurate than humans at detecting sexual orientation from facial images*, 114 J. PERSONALITY & SOC. PSYCH. 246 (2018).

[51] Nicolas Baya-Laffite, Boris Beaude, and Jérémie Garrigues, *Deep learning to predict sexual orientation in the public space*, 5 RÉSEAUX 137, 137-38 (2018) (challenging the Wang-Kosinski result).

widespread use.[52] By itself, at least, the carefully-cabined use of such technology to heighten the efficiency of building access creates no change in the quality or sheer quantity of state epistemic power.

Machine learning also helps unlock genetic information. Whether concerning single alleles or drawing on population-wide database of whole genomes, use of enormous computing power and population data to discern patterns in genetic information will likely continue to heighten the state's interest in learning what might once have been elusive if not downright unobtainable without face-to-face questioning of individuals. Recent studies have used genome-wide complex trait analysis and the use of polygenic scores (sometimes referred to as polygenic risk scores) to make predictions of social and perhaps political preferences.[53] Not surprisingly, scholars intensely debate the accuracy and implications of this work . On the one hand, a leading study from 2012 has pointed out that at least some genetic studies result at present in predictions that are only weakly powered.[54] On the other hand, a number of other studies find that single genetic traits related to serotonin transport provide a level of predictive acuity in respect to voter turnout behavior.[55] These incremental steps are unlikely to be the end of debates about the epistemic gains from genetic data.[56] At minimum, it seems not implausible that a piece of physical evidence (such as blood or spit) will at some point be able to reveal not just identity but also a range of preferences and behavioral traits.[57]

Finally, machine learning is likely to influence the power-related definition of privacy. Machine learning's availability shifts the "balance of power" between institutions and individuals. "[L]arge institutions—both governments and corporations—are gaining the upper hand … by tracking vast quantities of information about mundane aspects of our lives."[58] The acquisition of these large pools of data is of limited importance without machine learning, which provides the computational pathways to extract individualized information from them. Such tools, however, require large amounts of

---

[52] Equally, the use of translation software depending on machine learning tools, such as Google Translate, does not by itself appear to raise privacy concerns. Such examples show that an analysis should focus on the specific uses of a technology, rather than the technology in the abstract. But see Eugene Volokh, *The Mechanisms of the Slippery Slope*, 116 HARV. L. REV. 1026 (2003)(offering analytically-plausible rationales associated with multi-peaked preferences, attitude change, small-change tolerance, and political power and momentum to explain how policy or technological changes in one context may spread to other domains despite judicial efforts to cabin the process).

53 *See generally* David B. Braudt, *Sociogenomics in the 21st century: An introduction to the history and potential of genetically informed social science*, 12 SOCIOLOGY COMPASS -- (2018) (surveying recent advances in sociogenomics).

54 Daniel J. Benjamin et al., *The genetic architecture of economic and political preferences*, 109 PROC. NAT'L ACADEMY OF SCI. 8026, 8026 (2012).

55 Kristen Diane Deppe et al., *Candidate genes and voter turnout: Further evidence on the role of 5-HTTLPR*, 107 AM. POL. SCI. REV. 375 (2013) (replicating earlier studies, and adding new studies, to show a measure of predictive power); see also Sven Oskarsson et al., *Linking genes and political orientations: Testing the cognitive ability as mediator hypothesis*, 36 POL. PSYCH. 649 (2015).

[56] For a useful and succinct treatment of the perils of over-interpreting polygenic risk scores, see Michelle Meyer, Patrick Turley, and Daniel J. Benjamin, *Response to Charles Murray on Polygenic Scores* (Feb. 3, 2020), https://medium.com/@michellenmeyer/response-to-charles-murray-on-polygenic-scores-e768cf145cc.

[57] It is interesting to note that Justice Gorsuch discussed governmental acquisition of DNA from third parties as something that would violate reasonable expectations of privacy, notwithstanding the absence of genetic profiling in the case there at bar. *Carpenter v. United States*, 138 S. Ct. 2206, 2262–63 (2018) (Gorsuch, J., dissenting).

[58] JULIA ANGWIN, DRAGNET NATION: A QUEST FOR PRIVACY, SECURITY, AND FREEDOM IN A WORLD OF RELENTLESS SURVEILLANCE 29 (2014).

money, expertise, and computational power to employ.[59] Their use to acquire private information or to impinge on individual dignity is largely confined to the state and entities with the same level of resources as the state. At the same time, machine learning algorithms themselves can be opaque because they are not "explainable in human language,"[60] or else are commonly shielded from public scrutiny by legal regimes such as trade secrets.[61] Hence, even as they increase the capacity of the state to acquire information about its subjects, the instruments through which such acquisition occurs may become more difficult (or more costly) to understand. Just as public and private life becomes a degree more transparent to the state, so the state as a consequence of the same technology becomes less transparent to its publics.[62]

Asymmetrical access to information, as the drafters of the Fourth Amendment were aware, can be a potent instrument of political control. The Chinese state, for example, uses "facial recognition and artificial intelligence to identify and track 1.4 billion people" and thereby "assemble a vast and unprecedented national surveillance system."[63] Chinese police stationed at transportation hubs such as train stations, for instance, already use dark glasses with embedded data streams employing facial recognition technology, such that merely looking at a person pulls up their identity and related information.[64] An algorithmic classification tool sorts surveillance data for ethnic Uighur faces, producing a detailed accounting of the precise movements and actions of a single ethnic class.[65] Whether or not this technology works well—the Chinese government has not rushed to say—it may well be that technological change, along with improvements in the acquisition and cleaning of data will make such instruments meaningfully effective within the next couple of decades. One of the predictions to emerge from our analysis later in the paper is that even if the full effects of machine learning on political power have not materialized in consolidated democracies yet, it may well only be a matter of time because uses now associated with authoritarian regimes turn up closer to home.

Most of the privacy-relevant uses of machine learning we have canvassed here concern the extraction of private information from large pools of data. It is worth noting the possibility of another use that does not impinge directly on privacy values, but that does at least touch on related normative value—i.e., the individualized prediction of criminal behavior. To illustrate, imagine that the government employs a machine learning tool to analyze historical training data as a means of determining how money launderers can be identified. The machine learning tool predicts that a person

---

[59] This seems to us generally true, but there is some evidence that it is changing. *See* Cade Metz, *Good News: A.I. Is Getting Cheaper. That's Also Bad News*, N.Y. TIMES, Feb. 20, 2018, https://www.nytimes.com/2018/02/20/technology/artificial-intelligence-risks.html?action=click&module=RelatedCoverage&pgtype=Article&region=Footer.

[60] Tal Z. Zarsky, *Transparent Predictions*, 2013 U. ILL. L. REV. 1503, 1568 (2013). On the risk of opacity, see Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 25 (2014) (arguing for oversight and transparency).

[61] Rebecca Wexler*, Life, Liberty, and Trade Secrets: Intellectual Property in the Criminal Justice System*, 70 STAN. L. REV. 1343 , 1350 (2018) (documenting "the introduction of trade secret evidence into criminal cases"); FRANK PASQUALE, THE BLACK BOX SOCIETY 12-15 (2013).

[62] It is possible to design algorithms so that they are more 'explainable.' Deven R. Desai & Joshua A. Kroll, *Trust but Verify: A Guide to Algorithms and the Law*, 31 HARV. J.L. & TECH. 1, 10-11(2017). But notice that where the state is concerned, it will be the state's decision whether or not to do so.

[63] Paul Mozur, *Inside China's Dystopian Dreams: A.I., Shame, and Lots of Cameras*, N.Y. TIMES, July 8, 2018.

[64] Paul Mozur, *Looking Through the Eyes of China's Surveillance State*, N.Y. TIMES, July 16, 2018.

[65] Paul Mozur, *One Month, 500,000 Face Scans: How China Is Using A.I. to Profile a Minority*, N.Y. TIMES, April 14, 2019.

who has visited a casino in the last 90 days is more than not likely to be engaging in money laundering. The Government in response enacts regulation requiring those who visit casinos to report this fact. The resulting disclosures are then used to direct investigative resources related to fraud. To be sure, this simple hypothetical is vulnerable to the objection that criminals would quickly adapt (such that fraudsters would learn to wait 92 days before gambling with their projects, or else might migrate to unregistered online sites). But it still seems likely that regulated actors' behavior will not be perfectly elastic to the incentives created by the law, whether as a result of ignorance, inattention, or irrationality. The same exercise, moreover, can be framed with criminal actions that ought to be relatively inelastic, such as ideologically motivated violence[66] or child sexual abuse.[67] This species of 'individualized crime prediction' does not directly raise an informational privacy issue. Depending on how it is implemented, however, it might conceivably impose dignitary harms (e.g., requiring people to reveal a shameful aspect of their character, such as a gambling habit), or induce people to forego activities they would otherwise have pursued. As such, it would mark a new vector of state control over individual behavior. We think that such examples presently lie on the periphery of the privacy-machine learning interaction. But they are hardly trivial as a practical matter or implausible in the near future.

In brief, then, the deployment of machine learning technologies at scale in the last decade has already had, and will likely increasingly have, an impact on several elements of privacy. Given the state's ready access to growing computing capacity and at least the potential to gather enormous data—whether initially obtained by private actors, or directly through state surveillance techniques—government officials can more cheaply and easily learn more about who is doing what, with whom, and with what apparent intentions. Where the state is concerned, less material or digital information will be required to make an increasing number of inferences about a person, or alternatively to impinge on a private domain. The net result will be a widening of the power gap between the state and its individual subjects.

C. Technological Protections for Privacy

Yet the tango of privacy and technology moves in more than one direction. At the same time that we underscore the privacy impacts associated with the state's use of machine learning, it is also important to see that the same (or similar) technologies can be designed to provide a measure of protection from the intrusions on privacy enumerated above. Two examples warrant particular attention here.

First, a method has been designed for anonymizing datasets such that any query run on the dataset will produce the same result where a specific subject is included or excluded. This technique, which is known as "differential privacy," provides at least one technical constraint on the de-anonymization of at least some large data sets.[68] Differential privacy works by "deliberately add[ing] npise to computations, in a way that promises that any one person's data cannot be reverse engineered

---

[66] Existing models of the individual turn to terrorism, however, are deeply flawed. Aziz Z. Huq, *Modeling Terrorist Radicalization*, 2 DUKE F. FOR L. & SOC. CHANGE 39 (2010).

[67] For a skeptical account of machine learning-based predictions in one experimental context, see Philip Gillingham, *Predictive risk modelling to prevent child maltreatment and other adverse outcomes for service users: Inside the 'black box' of machine learning*, 46 BRIT. J. SOC. WORK 1044 (2015) (analyzing New Zealand's pilot program).

[68] Cynthia Dwork, *A Firm Foundation for Private Data Analysis*, COMM. ASS'N FOR COMPUTING MACHINERY, Jan. 2011, at 86, 91 (defining differential privacy).

from the results."[69] It is necessarily "built and reasoned about on a case-by-case basis,"[70] and cannot be used in all machine learning contexts. For instance, differential privacy does not provide a shield against the facial recognition technologies described above. At the same time, in its sphere of application, it can be quite powerful. Hence, a 2012 study suggested that Facebook's publicly released ad-related data has been subject to a number of modifications that enable a measure of effectual differential privacy.[71] Both Apple and Google have since then employed differential privacy for iPhone and browser metadata respectively.[72] The current consensus view among experts of differential privacy thus appears to accept its utility in respect to some kinds of data and under some circumstances, but to resist the idea that it is some kind of global privacy solution.

Second, innovations in the healthcare domain are instructive. Concerns about patient privacy have pushed some research institutions to develop "synthetic" datasets for analysis, rather than relying on real datasets that are amenable to reidentification.[73] In this method, existing data is used to construct a simulated dataset that has the same properties. The method is not absolute proof against reidentification. But a recent study concluded that the trade-off achieved through synthetic datasets between "the possibility of *measurably* small privacy leakage in exchange for perhaps mathematically provable protection against reidentification," tended to increase privacy.[74] Like differential privacy, therefore, synthetic data only works in some domains, and even then does not work perfectly. It is a 'leaky' solution, that improves but does not wholly mitigate the privacy dilemma (when it is applicable). At the same time, though, just as we should be alert to the possibility that intrusive tools will increase in efficacy over time, so too should we be aware that countermeasures such as differential privacy or synthetic privacy, will also improve over time.

At a very minimum, therefore, machine learning technology need not always and necessarily be viewed as privacy's nemesis. The existence of such counter-measures again suggests that how a technology such as machine learning impacts privacy depends on the circumstances of its adoption, and the technical choices embedded in its code. Differential privacy and a reliance on synthetic data are not discrete measures that individual users can deploy.[75] Rather, they are measures that are adopted, if at all, at an institutional level as part of the overall strategy of integrating machine learning into the

---

[69] MICHAEL KEARNS AND AARON ROTH, THE ETHICAL ALGORITHM: THE SCIENCE OF SOCIALLY AWARE ALGORITHM DESIGN 27 (2019).

[70] Arvind Narayanan & Vitaly Shmatikov, *Privacy and Security Myths and Fallacies of "Personally Identifiable Information*," COMM. ASS'N FOR COMPUTING MACHINERY, June 2010, at 24, 26, available at http://www.cs.utexas.edu/~shmat/shmat_cacm10.pdf. Differential privacy is unavailable, for example, if the learning algorithm draws data from a continuous distribution. Kamalika Chaudhuri & Daniel Hsu, *Sample Complexity Bounds for Differentially Private Learning,* 19 JMLR: WORKSHOP & CONF. PROC. 155, 155-56 (2011). Another critique of differential privacy focuses on the difficulty of knowing when and where it will be needed. Kate Crawford & Jason Schultz, *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*, 55 B.C. L. REV. 93, 99 (2014).

[71] Andrew Chin & Anne Klinefelter, *Differential Privacy As A Response to the Reidentification Threat: The Facebook Advertiser Case Study*, 90 N.C. L. REV. 1417, 1455 (2012).

[72] KEARNS AND ROTH, *supra* note --, at 47

[73] *See, e.g.,* Neha Patki et al., *The Synthetic Data Vault*, 2016 IEEE INT'L CONF. DATA SCI. & ADVANCED ANALYTICS 399, 400-10 (demonstrating a technique--the synthetic data vault--used to create synthetic data from five publicly available datasets).

[74] Steven M. Bellovin et. al., *Privacy and Synthetic Datasets*, 22 STAN. TECH. L. REV. 1, 50 (2019).

[75] Pursued at the individual level, privacy is a "luxury good" that requires "time, money, and technological expertise" to implement." ANGWIN, *supra* note 58, at 229.

performance of a policy function. Their availability will depend on the political economy of the institutions doing the adoption. It is that topic to which we now turn.

II. A Political Economy of Machine Learning, Privacy, and the State

A. Surveillance in a Two-Level Domestic/International Game

The way a computational tool is employed will depend on the opportunities provided by the social and institutional landscape. Artificial intelligence researchers refer to this in terms of "affordances."[76] We provide in this part a general account of the way in which exogenous social and political forces, both domestic and transnational, shape the relevant affordances of machine learning. In so doing, we take the methodological step of presupposing that a technology such as machine learning is not somehow 'self-applying.' It has no natural or inevitable patterning of uses in the world. Rather, its adoption and dissemination is a function of conscious choices, or, at worst, a negligent drift in institutional formation as a result of inattention and an absence of oversight.

To understand the forces that shape machine learning's privacy-relevant adoption, we adopt (with some modifications) the two-level framework famously developed by Robert Putnam in the international affairs domain to model the production of international agreements. In Putnam's rightly influential account, the negotiation of such agreements by chief executives occurs at two levels simultaneously--at the level of international diplomacy, and also at the level of domestic politics. Each leader sits at "the international table [with] his foreign counterparts," while at "the domestic table behind him sit party and parliamentary figures, spokespersons for domestic agencies, representatives of key interest groups, and the leader's own political advisors."[77] For an agreement to be secured at the international level, the chief executive must also satisfy the demands of those at the domestic table. Putnam's work suggested that a domestic constraint can either help parties define the contours of a potential agreement at the international level because it limits the chief executive's ability to make concessions during negotiations, or it may doom the agreement because it illuminates wholly the existence of a "win set," i.e., the domain of outcomes acceptable to the relevant domestic interest groups.[78]

The core insight we take from Putnam's work is the possibility of interaction between domestic and international levels of policy-making. Our analysis, however, does not focus on a negotiated output at the international level, nor upon the dynamics of negotiation per se. Rather, we focus on decisions about the use of privacy-salient machine learning tools as a surveillance technology at the domestic level. But a basic intuition of Putnam's two-level theory can be applied, mutatis mutanda, here as well: That is, the circumstances and forms of domestic adoption of the technology will be a function not only of a domestic 'game,' but also of an international 'game.' Moving beyond Putnam's original model, we think that the latter may facilitate policy change at the domestic level. Alternatively, the international game may hinder or even preclude the possibility of domestic policy

---

[76] Thomas E. Horton, Arpan Chakraborty, and Robert St. Amant, *Affordances for Robots: A Brief Survey,* 3 AVANT 70, 73 (2012) (discussing the use of the theory of affordances in the field of artificial technology in order to "develop better agents").

[77] Putnam, *supra* note 11, at 434.

[78] *Id.* at 433-51; *accord* Keisuke Iida, *When and How Do Domestic Constraints Matter?*, 37 J. CONFLICT RESOL. 403, 403-05 (1993)].

stability. In either case, domestic policy is a function of a complex interaction between domestic and international dynamics, which simultaneously impinge on a state's policy choices.[79] We think this dynamic is not unique to machine learning. It occurs with many other technologies. For instance, the global diffusion of critical internet resources, such as the domain-name system, means that when domestic interest groups lobby to secure online protection for their intellectual property, the resulting legislative efforts necessarily interact with global internet governance frames.[80] But we think that each technology will generate its own distinct game-form. The manner in which domestic-international dynamics shape key internet-design decisions will differ from the way in which they influence the diffusion of CRISPR-Cas9.

Our aim here is not to offer a precise prediction of how the two-level game will play out in respect to machine learning. Rather, it is to demonstrate that there is a two-level game in the first instance. As a result, any analysis of or prescription for the machine learning state and its surveillance capabilities must account for the two-level nature of its dynamics, and explain why it will prove stable under pressures from both levels. Given the preliminary and theoretical nature of our analysis, we sketch in general terms the way the two-level dynamic will likely unfold, rather than offering a misleadingly precise point-estimate prediction.

To that end, the following analysis focuses on demonstrating that a two-level dynamic exists specifically with respect to the adoption of machine learning tools. At the domestic level, the state decides on whether to adopt such tools in the context of their rapid private adoption, along with the rapid and extensive creation of deep pools of data necessary to exploit them. In the international sphere, democracies such as the United States must account for the deployments of machine learning tools for both internal and external purposes by other powers. In this domain, we think it is especially useful to focus on China as a consequential actor.

B. The Domestic Level

We begin by thinking about the ecology of interest groups that will influence the government's decisions to adopt, or to limit, machine learning tools with privacy-relevant effects. There is a robust coalition of interest groups who depend upon the information made available by machine learning tools, even where it can yield disclosures of private information or embarrassing slights to dignity. In a recent scholarly article and popular book, Shoshana Zuboff has dubbed this formation of companies, academic institutions, and investors a distinctive social formation called "surveillance capitalism" that comprises a "new form of information capitalism [that] aims to predict and modify human behavior as a means to produce revenue and market control."[81] The largest actors in this ecosystem are companies such as Google, Facebook, and Amazon.[82] According to Zuboff, all these entities follow

---

[79] Our approach is consistent with, although different in emphasis from the "New Interdependence Approach," which emphasizes instead the ways in which "globalization opens up political channels for other actors beside the state to engage in international politics." HENRY FARRELL AND ABRAHAM L. NEWMAN, OF PRIVACY AND POWER: THE TRANSATLATNIC STRUGGLE OVER FREEDOM AND SECURITY 3, 26 (2019).

[80] LAURA DENARDIS, THE GLOBAL WAR FOR INTERNET GOVERNANCE 2-6, 187 (2014).

[81] Shoshana Zuboff, *Big Other: Surveillance Capitalism and the Prospects of an Information Civilization*, 20 J. INFO. TECH. 75, 75 (2015).

[82] *Cf.* SCOTT GALLOWAY, THE FOUR: THE HIDDEN DNA OF AMAZON, APPLE, FACEBOOK, AND GOOGLE 1-12 (2017) (detailing the positions of these players, inter alia, in the surveillance economy).

the same three-part business model: acquire personal data inadvertently produced through interaction with digital goods and services; develope predictive models from these large pools of interaction-derived data to seek to advertisers; and then launch strategies for "behavioral modification" in order to make predictions even more attractive to advertisers (and so more lucrative).[83] In addition to the entities Zuboff stresses, there is a class of data brokers, or "companies that sell or exchange personal data" such as Experian, Axciom, Rapleaf and Datalogix, that comprise a roughly $200 billion industry.[84] The information held by those entities on specific individuals can be extensive, including all addresses and phone numbers used during adult lives, all relatives, every email contact and web search made, an account of shopping habits, and internal communications with employers.[85]

Although we have elsewhere raised concerns about elements of Zuboff's account,[86] we think the term surveillance capitalists is useful here to characterize those private entities that are likely to resist regulation of the private sector. The available evidence suggests that they are likely to be successful in that regard. In her account of Google and Facebook, Zuboff argues that both companies' chief executives have repeatedly shown "contempt for law and regulation" because their financial success depends on them "ignoring, evading, contesting, reshaping, or otherwise vanquishing laws that threaten [their supply of behavioral data]."[87] Even if one does not completely accept Zuboff's characterization of Google's and Facebook's attitude to the law,[88] it is hard to characterize Google's expenditure of $21 million and Facebook's expenditure of $13 million on federal lobbying in 2018 as an effort to go along quietly with the regulatory flow.[89] These actors, therefore, are a well-organized and powerful interest group. And where their lobbying efforts fail and regulation ensues, those companies also vigorously assert the First Amendment as a shield against legal constraint.[90] Adding to their power is the fact that government and surveillance economy firms are entangled in mutually beneficial dependencies. The latter are in effect national champions upon whom the country's economic success rides. At the same time, they are also vulnerable to regulatory hold-ups by the government, which can curtail or even stop their operations. These dynamics mean that private-sector development of privacy-relevant forms of machine learning are likely to flourish relatively untended by regulation.

---

[83] SHOSHANA ZUBOFF, THE AGE OF SURVEILLANCE CAPITALISM: THE FIGHT FOR A HUMAN FUTURE AT THE NEW FRONTIER OF POWER 512-16 (2018) [hereinafter "ZUBOFF, AGE OF SURVEILLANCE CAPITALISM"] (summarizing argument made in book as a whole); *see also* SCHNEIER, *supra* note 43, at 55 (noting the centrality of advertising to surveillance capitalism).

[84] Matthew Crain, *The limits of transparency: Data brokers and commodification*, 20 NEW MEDIA & SOC. 88, 93 (2018).

[85] ANGWIN, *supra* note 58, at 94-95.

[86] Cuéllar and Huq, *supra* note --, at 1309-25.

[87] ZUBOFF, AGE OF SURVEILLANCE CAPITALISM, *supra* note 83, at 105.

[88] By contrast, the record of large companies faced with government demands for data is decidedly mixed. *See Cooperation or Resistance?: The Role of Tech Companies in Government Surveillance*, 131 HARV. L. REV. 1722, 1725 (2018). This complicates Zuboff's view of these companies as scofflaws across the board.

[89] Ben Brody, *Google, Facebook Set Lobbying Records as Tech Scrutiny Intensifies*, BLOOMBERG, Jan. 22, 2019, https://www.bloomberg.com/news/articles/2019-01-22/google-set-2018-lobbying-record-as-washington-techlash-expands.

[90] ZUBOFF, AGE OF SURVEILLANCE CAPITALISM, *supra* note 83, at 109-10.

The power of surveillance capitalists can also be attributed to the fact that to date, legal arrangements rooted in familiar doctrinal domains such as contract or consumer protection do not readily seem for many observers to serve as an effective check on this surveillance economy. The leading regulatory strategy with respect to privacy is as familiar as it is an effective catalyst for eye-rolls: it entails requiring individualized consent before the acquisition of information. It has not proved a success. Privacy scholars have argued that disclosures are not only "vague and general," but also "tend to conflate important distinctions between remembering users' preferences, creating predictive profiles that may also include other, inferred data, using those preferences for targeted marketing, and tracking users across multiple websites, devices, and locations."[91] A yet more profound challenge to consent-based privacy regimes is that consumers seem to place different values on that good depending on whether they were asked to consider how much money they would accept to disclose otherwise private information or how much they would pay to protect otherwise public information.[92] That is, the baseline distribution of information seems to shape judgments about privacy's values. Consumers also appear to have time-inconsistent preferences, in the sense that they are (perhaps irrationally) willing to accept low rewards now in exchange for the "possibility [of a] permanent negative annuity in the future."[93] A final factor is the relative success of surveillance capitalists in obtaining a favorable regulatory environment. In stark contrast to the narrow applicability of consumer privacy protections, those companies' statutory immunities have been construed broadly. For instance, "internet intermediaries" benefit under the Communications Decency Act from wide protection them from liability based on the speech they facilitate.[94] In short, whereas privacy protections tend to fail, protections for entities that benefit from harvesting private data tend to thrive.

The relative absence of regulatory constraint on surveillance capitalist firms redounds to the benefit of the government in the form of an expanded capacity to acquire private information through those private intermediaries. As currently practiced in the United States and many other countries, surveillance capitalism presupposes that large segments of the public find it so convenient, or even fun, to use third-party services that elicit their information.[95] At least until recently, the resulting shared pools of data categorically fell outside the scope of federal constitutional protection under the Fourth Amendment as a consequence of the third-party doctrine. Accordingly, the more successful players in the surveillance economy are in acquiring data, the more data its law-enforcement elements can acquire without the cost of even a warrant. The recent Supreme Court decision in in *Carpenter v. United States*, holding that acquisition of cell-site locational data from a suspect's telecommunications provider counted as a "search" under the Fourth Amendment,[96] may signal a change in that doctrinal rule. How

---

[91] Julie E. Cohen, *Turning Privacy Inside Out*, 20 THEORETICAL INQUIRIES L. 1, 6 (2019).

[92] Alessandro Acquisti, Leslie K. John, and George Loewenstein, *What is privacy worth?*, 42 J. LEG. STUD. 249, 249-51 (2013).

[93] Alessandro Acquisti et al., *The Economics of Privacy*, 54 J. ECON. LITERATURE 442, 442-43 (2016). For similar results, see Kirsten Martin, *Privacy Notices as Tabula Rasa: An Empirical Investigation into How Complying with a Privacy Notice Is Related to Meeting Privacy Expectations Online*, 34 J. PUB. POL'Y & MARKETING 210 (2015).

[94] Kate Klonick, *The New Governors: The People, Rules, and Processes Governing Online Speech*, 131 HARV. L. REV. 1598, 1604 (2018) (discussing the judicial construction of §230 of the Communications Decency Act).

[95] SCHNEIER, *supra* note 43, at 58.

[96] 138 S. Ct. 2206, 2214–15 (2018) (considering whether acquisition of cell-site locational data from a third-party data provider constitutes a "search").

great a shift, though, remains to be determined through future jurisprudence.[97] And it is still possible that an opposite evolutionary dynamic could set emerge. That is, the federal courts might find that popular habituation to the pervasive sharing of data has the effect of changing what state activities impinge upon a reasonable expectation of privacy for the purposes of the Fourth Amendment. Habit might rub away at the reach of the Fourth Amendment, or at least work as a frictional constraint on its expansion.

Under the present doctrinal framework, the security elements of the government have strong incentives not just to allow surveillance capitalism to flourish, but to benefit from its epistemic fruit. Data is not just "an essential basis for economic exchange" but also "a potent source of control for government."[98] As legal scholar Jon Michaels has documented, "informal intelligence agreements with corporation" are already favored because they allow agencies to "direct broad swaths of intelligence policy without having to seek ex ante authorization or submit to meaningful oversight."[99] The result, he notes, is that "the intelligence agencies [already] depend greatly on private actors for information gathering."[100] Machine learning's advent exacerbates and accelerates that dependency.

Could pressure emerge to counterbalance these forces? Could, for example, inadvertant leaks of private data alter public perceptions? To date, though, breaches of data entrusted to by either the government or private entities have not proven focal points for public ire sufficient to have a material, longer-term impact on politics. But this could change. There may come a point when public frustration with the absence of data integrity, or the monetization of data by large technology companies at a time of income inequality, creates an opening for political entrepreneurs. A sufficiently large data breach might create a tipping point in privacy expectations and behavior.[101] Even without an exogenous shock, new norms of commercial actors competing for market share might generate new expectations of privacy that courts may be willing to recognize as 'reasonable.' A pro-privacy coalition might be the basis of a new political alignment, or at least a novel coalition. Or firms might respond to the existence of overlapping national regulatory regimes by seeking a more privacy-friedly compromise that the one struck domestically.[102] We should also not reject out of hand the possibility that courts might play a role in creating new focal points for privacy-oriented concern, or as providing a spur to legislative action at either the state or the federal level. Should such a realignment of preferences emerges in some form or another, it would create pressure not just on surveillance capitalists but only upon

---

[97] One constraint on the development of such principles is the *Leon* rule, pursuant to which certain good faith efforts of constitutional law will not trigger the exclusionary rule. United States v. Leon, 468 U.S. 897, 923–24 (1984). The effect of *Leon* and its progeny is to eliminate Fourth Amendment rightsholders' incentive to pursue their constitutional interests in the class of cases in which the law is least clear and most in need of clarification. *See* Aziz Z. Huq & Genevieve Lakier, *Apparent Fault*, 131 HARV. L. REV. 1525, 1550-51 (2018) (describing the effect of *Leon* on doctrinal development); *see also* Orin S. Kerr, *Fourth Amendment Remedies and Development of the Law: A Comment on* Camreta v. Greene *and* Davis v. United States, 2010 CATO SUP. CT. REV. 237, 237–39 (discussing convergence in these two lines of Fourth Amendment remedies). Another question is whether privacy will override the property rights that surveillance capitalists and others assert in aggregated data and metadata.

[98] FARRELL & NEWMAN, *supra* note --, at 18.

[99] Jon D. Michaels, *All the President's Spies: Private-Public Intelligence Partnerships in the War on Terror*, 96 CAL. L. REV. 901, 904 (2008).

[100] *Id.* at 907.

[101] For a fascinating fictional account of this, see BRIAN K. VAUGHN AND MARCOS MARTIN, THE PRIVATE EYE (2015).

[102] FARRELL & NEWMAN, *supra* note --, at 26-28.

government in respect to deployments of machine learning that impinge on privacy (as well as uses, such as the prediction of individual crime that only peripherally touch on privacy). It is likely that the leaders of such a broad-based public movement would have to overcome enormous barriers to engage in effective collective action. (Imagine, for example, the difficulty of organizing users of a market-dominant search engine to go on "strike" for a few days in order to force the kind of bargain that would yield a data dividend or higher-quality services for users[103]).

Nevertheless, it is important not to treat these difficulties as a reason to assume that it's impossible for policy entrepreneurs and social activists to overcome the so-called 'privacy paradox' arising from consumers' inconsistent preference, and to generate, in the process, a measure of new political cleavages and alignments.[104] Perhaps slowly, public perceptions may change about the relationship between individual economic well-being and the generation of data. A harbinger of such change is recent talk of a "data dividend" in the political rhetoric of states such as California.[105] Skepticism about the societal benefits associated with the practices of large technology companies, moreover, can cut across the political spectrum.[106] And at least in countries like the United States with relatively independent judiciaries, state and federal courts may yet play a role in that process as they help the public coalesce around particular ideas of what constitutes a breach of trust from public officials.[107] Courts can play this function, for example, if they more clearly articulate the scope of privacy rights in doctrinal contexts ranging from federal constitutional disputes,[108] to state constitutional interpretation,[109] to tort suits assessing the nature of the harm suffered from data breaches.[110]

Against these prospects stand large "surveillance capitalist" technology companies, and to some extent even government agencies. All reap concentrated benefits from the status quo, while the costs to the public are generally widely dispersed.[111] Vast segments of the public have accepted entrenched business models built on the putatively free provision of services in exchange for data. These arrangements have a path-dependent quality, in part because consumers unhappy with these arrangements face collective action problems similar to those citizens face in coordinating behavior to pressure their governments. Political polarization and partisan distrust, at least in the United States, lowers the probability of legislative deals that might have otherwise seemed particularly feasible because privacy issues can scramble the familiar left-right political divide. All of which probably leaves

---

[103] *Cf.* JARON LANIER, TEN ARGUMENTS FOR DELETING YOUR SOCIAL MEDIA ACCOUNTS RIGHT NOW (2018) (developing a set of, um, ten arguments in favor of a permanent social-media strike).

[104] *See*, e.g., M. Ryan Calo, *The Drone as Privacy Catalyst*, 64 STAN. L. REV. ONLINE 29 (2011).

[105] *See* Kartikay Mehrotra, *California Governor Proposes Digital Dividend Aimed at Big Tech*, BLOOMBERG (Feb. 12, 2019).

[106] *See The New Center Takes on Big Tech*, PR NEWSWIRE (Nov 26, 2018).

[107] Mariano-Florentino Cuéllar, *From Doctrine to Safeguards in American Constitutional Democracy*, 65 UCLA L. REV. 1398 (2018).

[108] *United States v. Maynard*, 615 F.3d 544 (D.C. Cir. 2010), *affirmed sub nom. United States v. Jones*, 565 U.S. 400 (2012)

[109] *People v. Buza* (S223698)(April 2, 2018)(Cuéllar, J., dissenting).

[110] Lior Jacob Strahilevitz, *Reunifying Privacy Law*, 98 CALIF. L. REV. 2007 (2010)

[111] *See generally* MANCUR OLSON, THE LOGIC OF COLLECTIVE ACTION (1971).

governments in advanced capitalist economies with considerable ability to rely on machine learning to monitor the public, either directly or through large technology company intermediaries.[112]

In sum, at least in the absence of discontinuities in the nature of policy entrepreneurship or social movements, the domestic political environment is one in which private entities (Zuboff's surveillance capitalists) have the ability to harvest and monetize large volumes of personal data, and in which government benefits from their activities (notwithstanding occasional gestures of concern about individual privacy). Government and surveillance capitalists, moreover, are in a sufficiently entangled relationship of mutual dependency that the former is likely to have either direct or indirect access to machine learning tools or their fruit. We think that these forces make it more difficult for machine learning instruments, whether in private or public hands, to be constrained by law. While a countervailing public movement can be imagined, it would face daunting barriers to organization and effectual intervention. Together, these dynamics create conditions in which both powerful interest groups and the government have strong, convergent interests in maintaining privacy-salient uses of machine learning relatively lightly regulated. Countervailing domestic forces will not dobut continue recurring to litigation, advancing doctrinal arguments to help police distinctions between private sector and government access to data, for example, or to raise constitutional or statutory questions about the use of machine learning to automate governmental decisions using enormous data. Yet given the extent of convergent interests in amassing information and consumers' apparent acceptance of surveillance-driven business models in exchange for conveience, domestic opposition to expanded surveillance will likely remain fragile—though not necessarily doomed to political failure.

C. The International Level

Political failure is what former House Speaker Tip O'Neill steadfastly avoided, and "all politics is local" is perhaps the phrase most associated with his storied career.[113] Although domestic pressures and concentrated economic consequences unquestionably matter, the only way O'Neill's statement can be right — at least in the realm of privacy and machine learning — is if we understand "local" to encompass the geopolitical realities constraining even the most powerful states. As countries navigate both domestic and international pressures when setting trade policy, deciding on the size and allocation of military budgets, and managing migration, so (we contend) it is with privacy. Which is why the second level in which a national government's decision on whether or not to adopt privacy-relevant machine learning tools must be understood as international and geopolitical in scope. The national government's decisions on such technology will necessarily be made in light of the parallel decisions to adopt or not adopt such tools by the nation's geostrategic opponents. Domestic nonstate actors will also use the international sphere as a channel for lobbying and advocacy on behalf of desired policy choice. This international level thus interacts with, and complicates, the domestic dynamics described above in complex ways. We set forth the basic dynamics, and then turn to the likely interactions.

---

[112] There is nothing inevitable about this in principle; Europe has taken a different path. FARRELL & NEWMAN, *supra* note --, at 46-53. But cultural and institutional differences no doubt exert a powerful influence on the likelihood of different legal, policy, and political changes in the United States relative to other countries. *See*, e.g., Mariano-Florentino Cuéllar, *Administrative War*, 82 GEO. WASH. L. REV. 1343 (2014)(discussing institutional and ideological constraints shaping the American approach to administering war mobilization and adapting public law on the eve of, and during, World War II).

[113] *See* Charles P. Pierce, *Tip O'Neill's Idea That All Government is Local Is How Government Dies*, ESQUIRE (Jul. 17, 2015)

Inter-state competition over machine learning, as well as the larger class of artificial intelligence-related technologies, is multi-faceted and complex. We have already mentioned China's use of machine learning tools, and China looms large as the other potential "AI superpower."[114] But it is not alone. Russia, South Korea, and Israel have also acquired the skilled personnel, large reservoirs of data, and computational resources to compete globally, with nations like India lagging behind.[115] As President Vladimir Putin's pronounced, somewhat bombastically, "artificial intelligence is the future, not only for Russia, but for all humankind, [and w]hoever becomes the leader in this sphere will become the ruler of the world."[116] These efforts to acquire leadership in the artificial intelligence space are occurring contemporaneously with national efforts in countries where the government is attempting to exercise more extensive control over internet-based communications.[117] The two efforts are not unconnected. For instance, China has complemented its so-called "Great Firewall," with what the Canadian research group Citizen Lab called a "Great Cannon," an offensive, internet-based system that "hijacks traffic to (or presumably from) individual IP addresses" and "can arbitrarily replace unencrypted content," so as to "manipulates the traffic of "bystander" systems outside China, silently programming their browsers to create a massive [distributed denial of service] attacks."[118] The Great Cannon is derived from the same code used to run the Great Firewall, and relies on the same pool of computational tools as machine learning.

Machine learning is a geostrategic asset because it can be used to amplify existing military capacities, sharpen cyberwarfare ability, and generate new security-related instruments. For instance, machine learning can be used to make targeted email attacks on adversaries security services more efficient; to mimic voices or create audio files that facilitate unauthorized access or spread disinformation; or even to target other automated systems (think of a hijacking by a foreign adversary of self-driving cars).[119] Moreover, "the properties of efficiency, scalability, and exceeding human capabilities" means that highly effective cyberattacks will become "more typical."[120] As a result of these military affordances, an "arms race" in the "vigorous prevention and mitigation measures" seems

---

[114] For a prediction of Chinese and American duopoly, see KAI-FU LEE AI SUPERPOWERS: CHINA, SILICON VALLEY, AND THE NEW WORLD ORDER (2018)

[115] *See* Tom Simonite, *For Superpowers, Artificial Intelligence Fuels New Global Arms Race*, WIRED, September 8, 2017, https://www.wired.com/story/for-superpowers-artificial-intelligence-fuels- new-global-arms-race/; Julian E. Barnes and Josh Chin, *The New Arms Race in AI*, WALL ST. J., Mar. 2, 2018, https://www.wsj.com/articles/the-new-arms-race-in-ai-1520009261.

[116] James Vincent, *Putin says the nation that leads in AI 'will be the ruler of the world*,' THE VERGE, Sept. 4, 2017, https://www.theverge.com/2017/2019/2014/16251226/russia-ai- putin-rule-the-world.

[117] *See, e.g.,* Vindu Goel, *India Proposes Chinese-Style Internet Censorship Rules*, N.Y. TIMES, Feb. 14, 2019, https://www.nytimes.com/2019/02/14/technology/india-internet-censorship.html; Rebecca MacKinnon, China's "networked authoritarianism," 22 J. DEM. 32, 34 (2011).

[118] Bill Marczak, et al., *China's great cannon*, 1-2 (2015), https://s3.amazonaws.com/academia.edu.documents/37269796/Chinas_Great_Cannon.pdf?AWSAccessKeyId=AKIAIWOWYYGZ2Y53UL3A&Expires=1553573854&Signature=LCM%2FYfPqupjVB3Sjr0ftF1D1p2Q%3D&response-content-disposition=inline%3B%20filename%3DChina_s_Great_Cannon.pdf

[119] Miles Brundage et al., *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation* 20-21 (Feb. 2018), https://arxiv.org/ftp/arxiv/papers/1802/1802.07228.pdf.

[120] *Id.* at 21.

likely (if not extant).[121] Perhaps in recognition of this international dynamic, the White House recently issued an executive order on "maintaining American leadership" on artificial intelligence.[122] The fiscal implications of perceived international competition over machine learning for surveillance capitalists such as Google and Amazon, we think, are fairly straightforward.[123]

Perhaps the most interesting geostrategic competitor to the United States will be China, where development of national AI resources has been elevated to the level of a "megaproject," and become the object of sustained multiyear investment and planning since 2017.[124] The existence of a Chinese "party-industrial" complex in which national champions like Baidu, Alibaba, and Tencent are seamlessly integrated into (and even lead) major initiatives on artificial intelligence means that dual-use technologies can be rapidly identified and implemented.[125] It likely also means that domestic security functions of artificial intelligence can be integrated frictionlessly into the manufactured fabric of goods and internet-based services. Perhaps the most interesting potential development is a convergence in artificial intelligence and quantum computing occurring at Tsinghua University's National Key Laboratory of Intelligent Technologies and Systems.[126] Were this convergence to be fully realized, it could "accelerate the process of machine learning for which computing capabilities remain a bottleneck at present."[127] It might, in other words, be a gamechanger at the geostrategic level.

D. Domestic-International Interactions Around Machine Learning

The international dynamics described here directly shape the domestic policy environment. On rough first approximation, our suggestion here is that their interaction makes the effective regulation of machine learning technologies less rather than more likely, whether they are in private or public hands (a distinction that is likely, in any case, to become increasingly blurred over time). Correspondingly, the international level of the privacy/machine learning interaction is likely to increase the rate of privacy violations. In all, the interaction of the international and the domestic hence quickens settlement on an equilibrium characterized by extensive private and public adoption of machine learning tools capable of invading privacy, with little by way of redress for those whose

---

[121] *Id.* at 45; Michael Horowitz et al., *Strategic Competition in an Era of Artificial Intelligence* 10 (2018), https://www.cnas.org/publications/reports/strategic-competition-in-an-era-of-artificial-intelligence.

[122] Executive Order on Maintaining American Leadership in Artificial Intelligence, Feb 11, 2019, https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/. The order itself exhorts agencies to "consider AI as an agency R&D priority" but does little concrete to achieve this. Id. §4.

[123] To be sure, those companies also have offsetting reasons to seek the mitigation of international conflict. It seems likely that such firms have an interest in maintaining strong global supply chains, access to global markets (including markets for data), and access to foreign talent through a relatively flexible immigration system. At a retail level, firms may thus take different policy postures.

[124] Horowitz et al., *supra* note 121, at 12-13 (describing the December 2017, the Three-Year Action Plan to Promote the Development of New-Generation Artificial Intelligence Industry).

[125] *Id.*, at 12.

[126] Elsa B. Kania & John K. Costello, Qu*antum Hegemony? China's Ambitions and the Challenge to U.S. Innovation Leadership* 18 (2018), https://www.cnas.org/publications/reports/quantum-hegemony.

[127] *Id.*

interests are sapped. This prediction, to be clear, applies to machine learning: It is not meant to generalize to other, unrelated technologies caught in a two-level domestic-international game.

There are several salient pathways through which this dynamic operates. First, and probably most importantly, the arms-race quality of international competition over the development of machine learning means that no participating nation can afford to slacken or pause with innovation and implementation for fear of losing a strategic advantage.[128] The brute force of international competition, that is, minimizes the space for domestic pro-privacy innovation, such as differential privacy or synthetic data sets, or relegates it to geostrategically peripheral actors. Where government regulation (whether of the government's own use of machine learning to invade privacy, or the parallel and entangled activities of private actors) can be opposed not only because it undercuts profits, but also because it impinges on national security, it is less likely to be enacted. Consistent with this dynamic, the jurisdiction with the most privacy-leaning regulatory framework—the European Union—is also a relatively minor player in the geostrategic sphere.[129] Lacking as great a stake in that context, Europeans have a freer hand when it comes to the regulation of privacy-salient machine learning tools.

Second, the surveillance capitalists who comprise the most powerful domestic political lobby in the United States are also the target of foreign cyberattacks.[130] In responding to those hacks, surveillance capitalists have turned to the federal government for aid,[131] although often with disappointing results.[132] The threat of foreign cyberattack, which is perhaps most crisply exemplified in China's "Great Cannon," only deepens the relationship of mutual dependency between those companies and the government. At the same time, the federal government increasingly relies on those same companies for military applications (much to the chagrin of some of their employees)[133] and raw data.[134] Hence, both the defensive and offensive elements of geopolitical dynamics work to intricate the interests of the surveillance economy and the government.

Third, the pursuit of geostrategic interests requires that governments account for the possibility that both personnel and projects will migrate overseas in order to arbitrage differences in regulatory stringency. Imagine, for example, that the federal government imposed a moratorium on individualized predictive technologies focused on the identification of future criminality. A computer

---

[128] Nick Bostrom, *Strategic implications of openness in AI development*, 8 GLOBAL POL. 135, 139-40 (2017). Bostrom advocates a more open model of technological development.

[129] For a sobering analysis, see *What would happen if American left Europe to fend for itself?*, THE ECONOMIST (Mar., 3, 2019), https://www.economist.com/special-report/2019/03/14/what-would-happen-if-america-left-europe-to-fend-for-itself.

[130] *See, e.g.,* Kim Zetter, *Google Hack was Ultra-Sophisticated, New Details Show*, WIRED, Jan. 14, 2010, https://www.wired.com/2010/01/operation-aurora/. The 2018 hack of Facebook, however, has not been linked to a foreign nation. Mike Isaac and Sheila Frenkel, *Facebook Security Breach Exposes Data of 50 Million Users*, N.Y. TIMES, Sept. 28, 2018, https://www.nytimes.com/2018/09/28/technology/facebook-hack-data-breach.html.

[131] Kim Zetter, *Google Asks NSA to Help Security Network*, WIRED, Feb. 4, 2010, https://www.wired.com/2010/02/google-seeks-nsa-help/.

[132] Kristen E. Eichensehr, *Public-Private Cybersecurity*, 95 TEX. L. REV. 467, 494-99 (2017). Having documented the absence of effectual government aid, Eichensehr nevertheless reports that "the government has an incentive to cooperate, or at least maintain open lines of communication [with surveillance capitalists]." *Id.* at 502.

[133] *Id.* at 500; *see also* Alexia Fernández Campbell, *How tech employees are pushing Silicon Valley to put ethics before profit*, VOX, Oct. 18, 2018, https://www.vox.com/technology/2018/10/18/17989482/google-amazon-employee-ethics-contracts

[134] FARRELL & NEWMAN, *supra* note 12, at 18.

scientist interested in that field has every incentive to migrate to a jurisdiction that does allow testing, and perhaps even implementation, of those technologies. As a result, the regulation of a particular form of privacy-limiting machine learning may well have the effect of facilitating a foreign adversary's development and deployment of that technology. To the extent that research into a given application is federally funded, moreover, even the withdrawal of such monies may well have the same effect. The basic dynamic at issue here arises in other scientific fields, where the practice of pursuing experiments prohibited in one jurisdiction by finding laxer testing grounds is called "ethics dumping."[135]

That these international pressures constantly intersect with the domestic (or, in O'Neill's terms "local") sphere is the essential backdrop for understanding how privacy and machine learning will affect the state. To a first approximation, these intersecting forces largely serve to exacerbate the tendency for governments to enable and encourage private action that undermines privacy (and that can be shared with government to allow for state intrusions on privacy), and to underinvest in technologies that might mitigate the negative privacy-related effects of such technologies. It is plausible to think that at least in the medium term, the domestic-international interaction in this field will lock in an equilibrium in which privacy is assigned low priority (again, in relation to private or public action alike), and privacy-reducing innovations are more rather than less likely.

Conclusion: Implications of Privacy's Political Economy in a Machine learning Age

We began by underscoring our debt to Stephen Schulhofer's clear-eyed and close read of institutional particularity. We have pursued in this essay a parallel project aimed at starting to map the political economy dynamics at work in relation to the development of machine learning tools with privacy implications. At a very general level, we have articulated a two-level (international and domestic) dynamic that presses toward acquisition and use. Based on this political economy, we suggest that there are few incentives *not* to collect data, and few incentives not to innovate in respect to the machine learning tools needful to its exploitation. In contrast, doctrinal innovation is likely to lag, not least thanks to the obstacles confronting any effort at statutory reform on the one hand, and the 'good faith' exception to the Fourth Amendment's immunity on the other hand. Technology may well present law with a *fait accompli*.

Our aim here is not to praise the resulting regulatory equilibrium, or to render its outputs as inevitable and therefore acceptable. Instead, our main effort here has been to understand and describe some of the domestic and geostrategic forces likely to shape the practices and policies affecting privacy in states where machine learning is deployed at scale in the public and private sectors. We think that given the likely pressures for greater use of privacy-relevant machine learning in government enforcement and related contexts provide an opportunity — and for some members of the public, lawyers, judges, and policymakers, a need — to clarify the values that society is trading off when it endeavors to protect privacy. Efforts to do so will no doubt engage familiar debates that have been raging for years. But we think it is clear that the terrain has been substantially changed by the emergence of new machine learning tools. Their advent comes at a moment in which social norms about privacy are changing, where a measure of political pressure on "surveillance capitalism" is crystallizing, and where decisions must be made on the future infrastructure for public and private

[135] *Recent events highlight an unpleasant scientific practice: Ethics Dumping*, THE ECONOMIST, Jan. 31, 2019, https://www.economist.com/science-and-technology/2019/02/02/recent-events-highlight-an-unpleasant-scientific-practice-ethics-dumping (discussing the use of Crispr-CAS9 to modify a human embryo as an example).

surveillance.[136] (e.g., 5G-enabled IOS, the shift from IPv4 to IPv6 and so on) is expanding. To the extent these dynamics produce unexpected openings for greater pro-privacy regulation, they raises the questions of how such efforts would be enacted and what their substantive content would be.

As these questions merit further exploration, we simply flag here some preliminary responses to these puzzles. At a threshold matter, it seems tolerably clear that Fourth Amendment doctrine, particularly as inflected by the third-party doctrine, provides only a limited vehicle for addressing the concerns raised by privacy-relevant machine learning tool. Since the latter emerge in both the private and the public sector, and since their epistemic gains can be easily triaged across the public private divide, a doctrine that hinges on state action is a poor fit. One response may be to modify the doctrine to account for the blurring of public and private functions.[137] Another might be, as we suggested earlier, to expect development of legislative approaches accounting for some of these issues, particularly at the state level. Indeed, in June 2018 Governor Jerry Brown of California signed the California Consumer Privacy Act, which will from 2020 create a "right to opt out" of the sale of "personal information about the consumer to third parties."[138] State courts, which can elaborate their own constitutions' privacy protections perhaps without regard to the state action requirement of the Fourteenth Amendment, are another pathway.[139] Yet another possibility could be grounded in (state or federal) administrative law, which might be used as a basis for regulations or internal operating procedures to vindicate privacy-related concerns. Such standards could make it easier to achieve a degree of convergence about the relevant norms among civil society leaders, responsible public officials, and members of the public concerned about privacy — convergence that could facilitate greater policing of transgressions when they occur.[140]

The full menu of conceivable (if difficult) regulatory pathways, and the substance of desirable regulation remain to be fully explored. Our aim here is to lay the foundation for that exercise by elucidating the two-level political economy that will continue shaping how privacy-related developments will play out in countries with organizations capable of deploying machine at scale. There is little prospect for any legal adaptation to technological change in this area that could reasonably be described as decent or equitable without an understanding of the pressures affecting how organizations will use machine learning to learn about behavior and what they will tend to with that knowledge.

---

[136] *See, e.g.,* Allan Holmes, *5G Cellphone Infrastructure is Coming: Who Decides Where It Goes?*, N.Y. TIMES, Mar. 3, 2018, https://www.nytimes.com/2018/03/02/technology/5g-cellular-service.html

[137] One approach would be to advance the "instrumental" approach suggested by William Stuntz, wherein "constitutional limits on law enforcement [are] aimed at minimizing the sum of the costs of crime and the costs of crime prevention." William J. Stuntz, *Local Policing After the Terror*, 111 YALE L.J. 2137, 2145 (2002). Machine learning changes the cost of producing security given its privacy-related effects. Hence, it ought to prompt clarification or modification of the doctrine.

[138] Assembly Bill No. 375, Ch. 55, § 1798.120(a), "An act to add Title 1.81.5 (commencing with Section 1798.100) to Part 4 of Division 3 of the Civil Code, relating to privacy," June 28, 2018. For an extended discussion, see Aziz Z. Huq, *A Right to a Machine Decision*, – VA L. REV. – (forthcoming 2020), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3382521.

[139] *Cf.* JEFFREY S. SUTTON, 51 IMPERFECT SOLUTIONS: STATES AND THE MAKING OF AMERICAN CONSTITUTIONAL LAW 177, 189 (2018) (exploring more generally the possibilities of state constitutionalism in relation to privacy protections).

[140] Cuéllar, *supra* note 93, at 1413-1414.