ESSAYS AND SCHOLARSHIP

Probably Speech, Maybe Free: Toward a Probabilistic Understanding of Online Expression and Platform Governance

Considering the cost of applying a probablistic statistical framework to First Amendment questions on digital platforms

BY **MIKE ANANNY**
AUGUST 21, 2019

FREE SPEECH FUTURES
An essay series reimagining the First Amendment in the digital age

Introduction

We are, again, asking what the First Amendment means, where its protections live, and who should enforce them. As legal scholars have argued,[1] the First Amendment is not only a legal doctrine or theory of jurisprudence but also a broader, often invisible, cultural force with looser constitutional groundings than many legal experts would prefer. The First Amendment lives not only in courts and case law but also in "First Amendment institutions"[2] —e.g., journalism, universities, libraries, churches, schools—that, collectively, pragmatically, culturally, define free speech, its protections, and the interests it privileges or harms. Understanding how these institutions see the First Amendment means searching for often hidden ways of thinking about speech that they share and debate. In this essay I suggest that *probability* is one of these largely invisible logics that has serious consequences for how we think about free speech and the First Amendment.

To understand when and why speech is created, chilled, censored, or celebrated, we can look not only to unconstitutional regulations on speech but also to the social, cultural, and technological forces that are "central to public discourse and its infrastructure."[3] In arguing for a more socially relevant "institutional First Amendment," Schauer asks courts to think about the power and particularities of institutions: "What distinguishes categories like viewpoint discrimination, content regulation, public forum, and prior restraint from categories like universities, libraries, elections, and the press is that the former exist in the First Amendment but the latter exist in the world."[4] Such institutions are not just "legal subjects that depend on the courts to tell them what the law is," they are places that

"develop their *own* visions of what the First Amendment means, even if that vision is different from the one courts would choose themselves."[5]

Today, the meaning and force of the First Amendment play out in the new and often unstable technological infrastructures[6] and institutional spaces of social media platforms. Variously seen as media organizations, technology companies, public spheres, private realms, and commercial spaces, these platforms solicit, moderate, circulate, interpret, and rank speech of all kinds.[7] From Facebook and Google to Instagram and Twitter, they "host and organize user content for public circulation, without having produced or commissioned it. They don't make the content, but they make important choices about that content: what they will distribute and to whom, how they will connect users and broker their interactions, and what they will refuse."[8]

Although scholars and practitioners alike hotly debate how and why to regulate such platform speech and regularly call for new types of accountability, they are still largely governed through self-initiated and self-designed policies, norms, and technological infrastructures.[9] Platforms largely govern themselves. They encourage users to flag offensive or harmful speech that break the community standards they write. They privately convene and manage content moderators to monitor speech and apply company policies. And they use proprietary machine-learning algorithmic systems to prevent some speech from ever circulating in the first place.[10] Increasingly, though, such self-regulation is proving insufficient. Scholars, activists, and technologists alike are trying to understand why flagging, moderation, and machine learning only sometimes work; they are showing the risks and limitations of relying exclusively on platform self-regulation and looking to social justice and human rights movements for inspiration and models; they are lobbying publics and lawmakers to force conversations about what government regulation or antitrust breakups might look like; and they are questioning the very existence of such platforms, suggesting that they are simply too fast and too big to be effectively governed.[11] With increasing urgency, people from myriad institutional contexts are seeing platform speech regulation as inadequate and fundamentally at odds with the idea of collective self-governance through free expression.

If we accept Horwitz's and Schauer's broad, institutional views of the First Amendment and agree that platforms are key institutions where the First Amendment's meaning and power play out, then we need to understand how platforms understand free speech. More specifically, what *logics* drive their regulation of speech? What assumptions, judgments, standards, definitions, values, and practices underpin their understanding of speech, and how closely do these logics align with normative ideals of what free speech *could* or *should* be? I argue here that *probability* is one of these logics.

Before examining probability and speech platforms more closely, though, I want to connect two bodies of work from outside the law to show how to see probability as a logic underpinning an institutional—not just legal—vision of the First Amendment. The first is

neo-institutional sociology and its core insight that organizational power is explained not only by looking for logics *within* companies or professions but by examining the "loosely coupled arrays of standardized elements"[12] that cut across institutions and tie them together. Some forces work *across* organizational settings to make institutions seem stable, predictable, teachable, and worthy of investment. For example, to understand medicine, we should look at how medical schools, hospitals, insurance companies, and state regulators make and remake the languages, routines, commodities, and expertise that make and remake definitions of illness, disease, and prevention.[13] Understanding "medicine" means understanding the myriad "situated knowledges"[14] that make some types of prevention taken for granted, some types of treatment controversial, and some types of expertise debatable. Similarly, to understand museums, we should trace how ideas about curation, preservation, funding, art, and professionalism emerge *across* the groups of artists, philanthropists, and audiences that define art, debate its stakes, train its practitioners, sustain its economics, and explain its value.[15] Scholars are similarly beginning to use these theories to trace neo-institutional definitions of journalism,[16] algorithms,[17] and press freedom[18] —looking for logics and forces that both contest and stabilize fields that live in no one place.

The second body of relevant scholarship, Science and Technology Studies (STS), asks how affiliations, knowledge, expertise, and risk play out in *technological* cultures. One of its key (and sometimes controversial) claims is that social life is understood not by measuring impacts of technology *on* society—artificially separating people and technology into causes and effects misses a fundamental point about how societies work—but, instead, by seeing social life as the inextricable intertwining of materials, practices, relationships, and values in infrastructures. The Facebook News Feed, for example, is not just a tool that ranks content. It is a technology that combines Facebook's business model, advertisers' interests, algorithmic signals, user actions, and myriad subcultures of speech, privacy, and social association. To think that the News Feed has an "effect" on its users is to mistakenly hold both the News Feed and people static, missing the fact that they simultaneously reflect and shape each other. STS scholars are concerned with what emerges from relationships between and among humans and nonhumans. They see things like knowledge, expertise, society, and risk *emerging from* the assumptions, outcomes, categories, stabilizations, aspirations, resistances, and breakdowns that result when people and machines are tightly, inextricably bound.[19]

With these images of institutions and technologies we can return to the First Amendment and ask a more precise question: What loosely coupled arrays of institutionally situated sociotechnical elements govern online speech? Put differently, which intersections of people and machines define the conditions under which people express themselves, circulate speech, encounter ideas, and suffer abuse online? When and how do particular sociotechnical intersections *matter* to free speech? Scholars of an institutional First Amendment could look across the often invisible institutional relationships and sociotechnical logics governing free speech and ask: Are these relationships and conditions

creating the First Amendment that publics need? If not, should we change the sociotechnical logics, or should we revise our understanding of the First Amendment?

For the remainder of this essay, I want to show how these approaches might be used to understand online speech platforms by focusing on one sociotechnical logic: probability. Probabilistic ideas—about chance, likelihood, normalcy, deviance, confidence, thresholds—underpin many of the sociotechnical infrastructures and institutions that regulate online speech platforms. They are often "baked in" to platforms, residing in complex and opaque systems that are hard to see or understand. Further, probability offers a kind of false stability couched in mathematical certainty that is beyond the comprehension of most platform users and regulators (and some makers) but that is routinely offered to provide an illusion of normalcy and predictability. Probability is a type of power that platform makers have a vested interest in obfuscating, mystifying, and controlling.

Just as we understand medicine better by seeing *disease* as a contestable concept across hospitals, schools, and patients, and just as we appreciate how *art* emerges from the intertwined power of patronage, curation, and craft, we may better know what free speech means online by critically interrogating *probability*—as an institutional and sociotechnical phenomenon. If we can see the contingencies and contestations of platform probability, we may not only discover it to be a largely invisible form of platform governance, but we may also be able to make better interventions into platform cultures, saying more precisely why certain forms of probability succeed or fail in their treatment of online speech.

Defining Probability

To know how probability fails or succeeds as an institutional technology, we need to understand a bit about its history and varied meanings. The story of probability is the story of how mathematics gains social power, how counting technologies become political instruments, and how technologists of every era distribute risk and claim confidence at increasingly large and complex social scales.[20]

In their history of chance,[21] Gigerenzer et al. argue that the idea of probability came from the need to formalize expectations, reduce risk, and make visible, verifiable, and defensible judgments about the future. As populations grew, commercial networks expanded and human relationships strayed from face-to-face interactions into increasingly mediated settings. People realized how complicated it was to build trust over distances, to decide which strangers to trade with, and to know which risks were worth pursuing. As the usual sources of authority (religion, philosophy, local community) were called into question or became increasingly complex, people searched for other types of certainty. Influenced by emerging legal doctrines about what human behaviors could *reasonably* be expected, mathematicians developed a science of "probable knowledge" based on three sources of authority: "physical symmetry" (e.g., coins had two sides and thus offered 50/50 odds), "observed frequencies of events" (e.g., weather incidents could be counted and aggregated

into predictable climate patterns), and "degrees of subjective certainty or belief" (e.g., some types of evidence were more or less reliable when trying a legal case or petitioning a judge). [22] Gamblers, farmers, insurers, and lawyers, eager for systems of certainty, were all drawn to the new science of prediction that stabilized environments, making it easier to know how to act within them. Betting, planting, underwriting, and arguing could all move away from superstition and luck. Instead, they could outsource their fears and concerns to probabilistic technologies—dice, coins, frequency tables—that were backed by mathematical and bureaucratic authority, regularly audited, and cared for by trustworthy statisticians.

These systems of probability, though, were limited by their reliance on observed frequencies. They were fundamentally descriptive and only worked within specific stable scenarios. They had little to say about how to *make* the world be more predictable or how to understand what happened when different conditions intertwined and outcomes were not always easily observable. These limitations laid the groundwork for two innovations in probability: *control* and *conditionality*. Though a discussion of these forces may seem to take us away from a focus on free speech, they are arguably at the heart of how platforms understand and regulate mass-scale online expression.

Control

To enable *control* through probability, new and powerful mass-scale counting practices, technologies, and values began to emerge. Surveys, censuses, standardized observations, philosophies of positivism, and behavioral categories all showed how probability could do more than simply standardize expectations about the world. Probability could show which parts of the world were *unpredictable* and needed to be controlled. The discovery of standard distributions of behaviors meant that people were "normal if they conform[ed] to the central tendency of such laws, while those at the extremes [were] pathological."[23] Techniques of probability could say with certainty exactly how common or rare an extreme was and how its variance appeared across a population. If those extremes could be tied to people and behaviors, then they could justifiably be surveilled for deviance, punished for straying too far from expectations, and held up as examples of how deviances harm collectives. Individuals—who were too unpredictable and idiosyncratic for statistical certainty—were reshaped into *groups* that behaved predictably and rationally so that a new "science of social physics" could produce *l'homme moyen*—the average man—an entity that voted, purchased, moved, lived, and died in ways that could be seen, aggregated, and shaped. These observations, categories, and databases made societies not just observable but *governable*.[24] Probability became a technology for "making up people."[25]

But hints of resistance emerged to claim that people are not the same everywhere. In their history of nineteenth-century probability, Gigerenzer et al. describe how systems of probability in England, France, and Germany represented different human values.[26] English statisticians defended the idea of free will and cautioned against ascribing patterns to individuals: "probability statements should be interpreted as predictions of long-run

frequencies rather than as quantification of uncertainty in a particular case."[27] French statisticians and officials were similarly cautious about probability becoming a tool for controlling individuals but invested heavily in state statistical bureaucracies focused on the idea that understanding the probable behavior of *l'homme moyen* was the key to making good public policy. People may have free will, but probability could still be an efficient and responsible way to see patterns of autonomy and allocated public resources.[28] German statisticians were the most skeptical of probability's damage to free will. To the Germans, the "idea that society could be typified by an average man simply reflected an impoverished conception of the human community." The job of the German statistician was to "break a population up into its various parts . . . in order to learn something about causes" and their varied sources.[29]

Similarly, in his study of how "socialist statistics" in the USSR, India, and China emerged as academic fields, bureaucratic practices, and a symbol of self-knowledge and "hallmark of a modern nation-state," Ghosh finds that that the very idea of a "social fact" depends on the type of society you think probability serves.[30] In contrast to Western traditions that saw statistics as a mathematical science for reducing societal randomness through random sampling, regressions, and probability theory, Marxist (Soviet-Chinese) statistics saw itself as a social science that rejected the idea of randomness or chance. Instead, it stressed the importance of *completely* counting everything that could be counted, preparing regular and exhaustive reports of populations, and creating sampling demographics and contexts that could be defended as "typical."[31] If Western statisticians tried to construct the average person and predict her behavior,[32] Marxist mathematicians aimed to capture entire societies and explain their movements.

Conditionality

Motivated by a desire to study not just patterns but also their variance, some statisticians recast the field as a way of understanding the *conditions* that influence outcomes. Broadly grouped under the umbrella of Bayesian statistics, the key insight is that models of probability need ways to incorporate new evidence—to learn about the world and change in response to it. If "frequentists" are concerned with the chance of X happening, "Bayesians" try to understand the chance of X happening if Y or Z has also happened—or might happen.[33] It means that Bayesians are not content simply to predict the world as it is; they are more speculative and experimental, eager to understand the set of possible outcomes if we had evidence about *other* outcomes. Rather than asking "is someone 20 versus 40 years old more likely to commit a crime?"—a frequentist would count instances and arrive at an answer—a Bayesian reframe the question as "under what conditions is someone 20 versus 40 years old more likely to commit a crime?" This puts conditions into play and raises questions about which conditions might matter and why: Does a previous conviction matter? Does gender matter? Does the day of the week matter, or the weather?

Bayesians are building models of the world by observing how the world behaves under conditions, but such models should make us ask where these conditions come from, which conditions are okay to ask about, and which conditions are unacceptable bases for governing social systems. Those conditions can come from existing datasets that describe the world, experimenters' folk theories, and simulations that imagine complex realities that do not actually exist—if Y *were to happen* with frequency *f*, how confidently could we predict the likelihood of Z? With enough data, no hypothesis is untestable, no string of contingencies unexaminable.

This Bayesian approach maps well to the "big data" made possible by omnipresent surveillance technologies and cheap information storage as well as the computational experimentation enabled by machine learning techniques and multivariate analysis. Instead of crafting *particular* experiments—motivated by theory to create or sample conditions under which beliefs may be true with some degree of probability[34] —today's Bayesians can use massive data sets and computational power to ask which conditional beliefs meet acceptable thresholds. Bayesian models can probably find enough evidence to support any set of conditions and beliefs. Machine learning statisticians can set belief thresholds and ask: Under which conditions are these belief thresholds met? Instead of asking how likely is X if Y or Z happen, they can ask: Since we have enough data on X, Y, or Z happening (with new data arriving constantly), what is the *complete set of beliefs* that meet a particular threshold in any given moment? Given a dataset of instances, what are all the things I can believe, now, if I want a confidence level of 95 percent? 80 percent? 70 percent? And what *could* I believe in the future, knowing that new data are constantly arriving?

★ ★ ★

While a complete history of the idea of probability is beyond the scope of this essay, what begins to emerge is a complex landscape of what probability *could* mean and an image of how these different meanings represent political choices and technological power.

Does probability live in natural systems that have inherent likelihoods—like coins, dice, or the weather? Is probability about observing frequencies of events and building models of as many outcomes as can be imagined and seen? Is probability a control technology, a way of knowing what is likely, normal, deviant, or expected in order to decide how to invest governance and which categories to reward? Is probability actually about sustaining beliefs under particular conditions? To hold a particular belief, what *else* do we need to know, and how confidently do we need to know it? Or, in an age of large-scale data and near-instantaneous computation, is the very idea of probability an anachronism—a holdover from a time when we could not count completely, were forced to find ways to "tame chance,"[35] and could not yet implement Marxist ideals of complete enumeration for control of entire populations?

Platform Probabilities and Speech

It may seem as if I have strayed significantly from my focus on First Amendment institutions, free speech, and platform governance. My claim, though, is that if we can see platforms asprobabilistic constructions, we might better recognize and redirect the probabilistic logics they use to govern expression and control populations. We never step into the same Google Search or Facebook News Feed twice because the companies that make them are constantly changing their ranking and recommendation algorithms, continually A/B testing new interfaces and options on population samples, and updating the behavioral models they use to convince advertisers that they can reliably reach particular markets.[36] Most fundamentally, platforms *are* probabilities because there is only a chance of them existing in any particular form at any given moment.

Platform content moderation is also probabilistic. It is a confluence of likelihoods: did an algorithmic filter trigger a computational threshold to block offensive content, did enough users within a particular period of time flag a sufficient amount of content to cause an account to be suspended, and did third-party content moderators evenly apply platforms' content standards? Many users simply do not know how their content is being moderated,[37] much less the shifting statistical ground on which such judgments stand. As Facebook's Monika Bickert acknowledges, "A company that reviews a hundred thousand pieces of content per day and maintains a 99 percent accuracy rate may still have up to a thousand errors."[38]

Even the news headlines, ledes, and advertisements we read are probabilistic—the result of constant, invisible A/B testing as publishers learn which words and images drive traffic and attract demographics. Depending on how confidently platforms can detect them and how interested they are in banishing them, some platform speech—e.g., approximately two thirds of tweeted links[39] —comes from automated bots and fake accounts designed to flood online spaces and create the impression that a viewpoint *might* be true because it is expressed so frequently.[40]

Similarly, mis- and disinformation are probabilistic phenomena, as makers and detectors of "deep fake" media play continual cat-and-mouse games to create and catch fabricated images, audio, and videos.[41] When Facebook partnered with news and fact-checking organizations to remove such content from its platform, it did not *delete* such content from its site; rather, it celebrated its statistical ability to "rank those stories significantly lower" and "cut future views by more than 80%."[42] After criticism that it spreads conspiracy videos, YouTube announced that it would "begin *reducing* recommendations of borderline content and content that could misinform users in harmful ways."[43] The move only impacts "recommendations of what videos to watch, not whether a video is available on YouTube." To be clear, Facebook and YouTube do not say they are *unsure* whether content is false or conspiratorial; they use private partnerships and proprietary algorithms to categorize speech, label such categories with certainty, and then *use probability* to strike balance they have define as "maintaining a platform for free speech and living up to our responsibility to users."[44]

Financial markets, too, reflect probabilistic speech systems. Thomson Reuters and Dow Jones have created systems to parse economic news stories, decide how confidently they understand the story's meaning, and then drive nearly instantaneous algorithmic trading decisions based on which confidence measures have been met. Even the creators of computational systems that autonomously produce speech have publicly warned that automatically generated news stories pose significant risks to public discourse because they find that readers are too quick to trust algorithmically produced speech; they underestimate the likelihood of encountering algorithmically generated text and overestimate their ability to distinguish them from human-written stories.[45]

We only ever have a chance of encountering speech as others do right now, or as we did days or even minutes ago. Our experiences with online speech are the product of loosely coupled arrays of sociotechnical systems: people, algorithms, commodifications, thresholds, confidences, intentions. At any given moment, this array drives the probabilities that speech appears, circulates, is believed, has monetary value, and drives action.

To unpack this further, consider three domains in which probability governs speech: the chance that content is banned, platforms' relationships to the risk profiles of professions, and the extent to which technology companies deflect responsibility for probabilistic outcomes to technical infrastructures and third-party designers.

Probably Banned

First, the idea of a speech "ban" makes little sense in probabilistic online environments. Platforms only ever make it more or less *likely* that speech circulates; they never guarantee the distribution or disappearance of speech.[46] This likelihood depends on a mix of factors determining how content is: created by a blocked, muted, or hidden user; flagged by users; caught by moderators; understood by translators; classified by humans interpreting community standards; sensed by machine learning algorithms; judged similar to previously blocked content; and highlighted by mainstream media outlets.

In practice, whole languages and regions are never susceptible to such bans because they are not being algorithmically monitored. Though Facebook offers its interface in 111 languages, its algorithms can detect hate speech in just 30 languages and "terrorist propaganda" in 19 languages. Its director of public policy for Africa said that "a lot of people don't even know that there are community standards"—and thus fail to flag speech.[47] No platform has translated its standards and algorithmic tools into all of the languages its interface supports. Such bans are probabilistic—never binary—and only applicable in languages and regions that are monitored.

For example, after reports that the 2017 London Bridge attackers were radicalized, in part, through extremist YouTube videos, parent company Google said that videos it defined as offensive but not in breach of its community guidelines would carry warnings, not be recommended, and be ineligible to earn advertising revenue. The goal, Google General

Counsel Ken Walker said, was to make these videos "have less engagement and be harder to find." Additionally, although he said that Google's "Trusted Flagger reports are accurate over 90 per cent of the time," the program would be expanded to include an additional 50 organizations to be trusted flaggers.[48] The additional flaggers would presumably increase the 90 percent figure to an even more acceptable number. Here, calculation represents trustworthiness: mistakes had been happening 10 percent of the time, but the attacks made that old number unacceptable. It had to decrease by some amount.

In internal discussions about whether to ban Alex Jones from its Instagram platform, Facebook used similarly actuarial thinking to defend its initial inaction. The company's "risk-and-response team" found that Jones's account failed to meet its violation threshold: "an IG [Instagram] account has to have at least 30% of content violating at a given point in time as per our regular guidelines."[49] Note the snapshot nature of this threshold: at least 30% violating content *at a given point in time*, not 30% over the account's lifetime. The harm threshold is calculated instantaneously, not cumulatively, and stays the same regardless of context. A U.S.-based Facebook executive went on to say that even the comments on Jones's account were under an unspoken threshold of acceptability: "The 560 comments have been reviewed. Only 23 [4%] are violating and therefore the object does not meet the threshold for deletion."[50]

Whereas YouTube's 90 percent figure was meant to reassure users that its mistakes were rare *across* its platform—90 percent was good and not as bad as 80 percent, but still not good enough to stave off calls for reform—Facebook's numbers (30% of violating content, 4% of comments) were seemingly commonsense evidence that only a small part of a *single* account was offensive, not yet offensive enough to warrant action.

Percentages, anthropologist Jane Guyer argues, "are better seen as performatives, aimed to call forth a judgmental response about 'the way things are going': it's getting too much; it's not fair anymore; it's reaching a danger zone."[51] These judgmental responses are, of course, contingent (when is too much too much?) and contextual (who is 90% certainty or 30% violations good enough for?). In online contexts where bans are impossible—Facebook's head of artificial intelligence admits "it's never going to go to zero"[52] —it seems past time to debate normative questions about who has the power to define bans as probably "good enough" and for whom.

Professional Alignment

Probability also can be a way to earn legitimacy by aligning with professions that are already socially accepted as imperfect. Such partnerships can be a way of sharing risk, communicating a system's imperfection, and managing users' expectations.

For example, in 2017 Facebook began deploying "suicide prevention tools that use artificial intelligence to identify posts with language expressing suicidal thoughts."[53] It developed "FBLearner," a proprietary "machine learning engine, to train a classifier to recognize posts

that include keywords or phrases indicating thoughts of self-harm."[54] Putting aside legal concerns about such systems operating outside of "HIPPA privacy regulations, principles of medical ethics, or rules governing research on human subjects," the system's success rests upon several types of probability.[55]

Relying on medical practitioners' statistical models of how words and phrases correlate with suicidal thoughts and actions, the system translates those probabilities into an algorithmic detector. Probabilities from one domain (clinical evidence) are imported into another (social media posts). Posts are parsed, classified, and correlated with suicide indicators; if certain thresholds are triggered, the system alerts emergency responders, who perform "wellness checks" on the possibly suicidal people.

To be successful, the system has to reliably transfer and translate probabilities from one context to another; social media environments have to look enough like clinical settings that each other's watchwords can be compared confidently enough to trigger action. Designers and clinicians also have to consider the potential effects of such wellness check systems and their popularity: Will knowing that social media environments are being monitored for suicidal language make at-risk individuals more or less likely to express themselves in such places? If human observers know that such a system is in place, will it make them more or less likely to intervene when they see potentially suicidal language, or will they trust that an algorithm driven by clinically reliable patterns is making a better judgment than they would? And consider the system's tolerance of false positives and false negatives. If the system errs and sees a suicide risk where none exists, or where the risk does not need a standardized wellness check, will suicide-related speech travel elsewhere and stand a reduced chance of being detected when a wellness check might be appropriate? Or, if Facebook claims to be on guard for suicide-related language and then *fails* to see genuinely suicide-related risks, has it created false comfort and unreliable expectations for suicide victims and their potential support networks?

To be sure, no system—human or computational—for recognizing suicide is error-free, but this system rests on new types of probabilistic assumptions: that suicidal language in social media platforms means the same things as in other settings; that algorithms can confidently parse subtle language, humor, sarcasm, code switching, and complex social contexts; that people share expectations about what the system is, what it is capable of, and how it compares to human judgment; and that the system's failures and risks of false positives and negatives have been fully anticipated and understood. As with the previous example in which the Facebook AI head acknowledged the impossibility of completely banning content, Facebook similarly issues caveats regarding the riskiness of this system by saying that "we're not doctors, and we're not trying to make a mental health diagnosis."[56]

Infrastructured Probability

Finally, probability can be a way to deflect responsibility for the configuration of technical infrastructures. In its study of Amazon's face recognition tool *Rekognition*, the American

Civil Liberties Union found that the technology falsely—but confidently—matched the images of 28 members of the United States Congress with people who had previously been arrested for committing crimes.[57] In its defense of the technology, Amazon did not claim that the tool had erred; rather, it said that the ACLU had failed to apply the proper confidence threshold:

> "While 80% confidence is an acceptable threshold for photos of hot dogs, chairs, animals, or other social media use cases, it wouldn't be appropriate for identifying individuals with a reasonable level of certainty," the [Amazon] spokesperson said. "When using facial recognition for law enforcement activities, we guide customers to set a higher threshold of at least 95% or higher."[58]

In response, the ACLU noted that the tool's default "similarity threshold parameter" is set to recognize faces that are considered 80 percent similar, a default that Amazon's technical documentation says can be changed.[59] In this example, probability is cast as an institution's *technical* responsibility. The ACLU critiques Amazon for creating a tool with an unacceptably low default confidence threshold and not contractually requiring that different thresholds be used in particular contexts. In making it possible for law enforcement officials to license the system and use any standards of certainty they wish, Amazon adopts a standard "buyer beware"defense, suggesting that its technology is just a neutral tool and that it is up to clients set the appropriate confidence threshold. To the ACLU, probability is the system's linchpin of certainty where the power of prediction plays out; to Amazon, probability is just another option for its customers to set as they like. Responsibility and accountability plays out in the interpretation of probability.

Further examples abound, but what we can see across all of these examples is how probability is simultaneously an *institutional achievement* (multiple actors convene to debate what probability means and what thresholds of risk, error, or confidence are acceptable), a *sociotechnical construct* (probabilities live in algorithmic sensors, large scale datasets, approaches to machine learning), and a *normative defense* (mistakes will be made, technologists are not clinicians, clients know best). Probability is, at once, a seemingly neutral technique, evidence of power, and a rationalization of risk.

As both its histories and contemporary applications show, probability does not mean any one thing. It can be evidence of expectations about the world and *observations* of seemingly natural frequencies. It can reveal attempts to *control* the world through categories used to define normality and punish deviance. And it can mean *experimenting* on the world, discovering the conditions under which certain people—and machines—can confidently hold expectations about patterns and outcomes.

These differences emerge as statisticians debate their techniques and grapple with how the field should see its social and humanistic underpinnings. They also play out in lawmakers'

images of how to govern people, what a "reasonable" person can expect in likely scenarios, and what it means to administer justice when cases sit between the patterned and the particular. And probability appears in the tacit knowledge and work practices of technological cultures that make, test, and deploy statistical systems. As technologists anticipate, simulate, and normalize error, they calculate systems to be "good enough" to deploy into the wild. Deploying a technology also means releasing benefit and risk into the world—enrolling people in failures you know will probably happen, forcing them to live your probabilistic calculations, and relying upon them to report or in some way reflect your errors. Such an intertwining of risk and probability is not new. As Dryer shows in her insightful history of probability's role in the politics and design of algorithmic systems of the 1920s onward, probability has always encoded—in practices, artifacts, values—human desires for certainty, control, and objectively defensible risk conclusions.[60] Probabilities are not just mathematical constructs or even a social constructions; they are *diagnostics* of the errors that people and societies are willing, able, or forced to endure. Probability lets us ask precise questions about how consequences are imagined, hoped for, endured, and resisted. It is about making *public* life.

Probability and Free Speech

Probability matters to free speech and free speech platforms precisely because the probabilities governing communication environments shape our collective ability to see and understand unavoidably shared collective outcomes—to discover ourselves *as publics* and know our chances of self-governance.

If we take seriously the idea that the First Amendment as currently conceived runs the risk of becoming irrelevant[61] and accept that we need a more expansive vision of the First Amendment as an *institutional* phenomenon, then we need to know all of the ways that speech is elicited, chilled, and celebrated. Probability is one of those ways. Through their designs, business models, and moderation policies, platforms are constantly making probabilistic, actuarial calculations with the power to shape collective self-governance.

Scholars and practitioners alike are only just beginning to understand how platforms' sociotechnical dynamics reveal dramatically uneven distributions of probability. Although probabilistic errors may be modeled and represented as collective, shared consequences, they are experienced by particular populations and individual people. How are probabilities distributed? When Facebook says that its fact-checking infrastructure catches 80 percent of falsities on the platform, who endures the harms of the other 20 percent? Or when it says that Alex Jones's Instagram account isn't offensive enough to sanction, who endures the unsanctioned hate? When Google says that its flaggers are over 90 percent accurate, who lives with the 10 percent error? When Facebook's suicide detection engineers acknowledge that "there will be mistakes," how do we translate this defense into a more active voice—it is not that "mistakes will be made," it is *Facebook* making mistakes—and hold them accountable? Indeed, what is the right unit to hold accountable? An individual mental

health practitioner who repeatedly made mistakes would be sanctioned differently from a clinic that erred frequently, and differently again from an entire profession that failed predictably. Where exactly *is* probability and error—and the accountability thereof—in speech systems that intertwine users, designers, algorithms, regulators, and venture capitalists?

We might also consider how different geographic and cultural meanings of probability play out in speech systems. As discussed, different statistical ideologies and practices have appeared, in various ways, in France, England, Germany, India, China, the United States, and the USSR. To what extent do these variations reflect different cultural understandings of statistics, of the types of surveillance and data gathering required to observe and predict behavior, and of the types of error and confidence statistical systems are culturally thought to have? If statistics is not only a mathematical technique but also a social fact, should the machine learning systems designed to monitor and govern platform speech be designed differently before being deployed in particular regions, cultures, or languages? Platforms already fail to provide machine translation tools for all the languages they support, but should the underlying logics and values of those tools differ?

It may seem esoteric, but the quest for certainty may be harming the planet. Curious about the environmental costs of artificial intelligence, computer scientist Emma Strubell and her colleagues studied the carbon footprint of natural language processing (NLP) technologies, like those that have been used to write compelling "fake news" articles.[62] To be reliable enough for peer-review publications and enterprise-level technologies, such technologies require vast amounts of algorithmic training and computing power, quickly becoming energy-intensive processes with significant environmental consequences. They found that training such a model "can emit more than 626,000 pounds of carbon dioxide equivalent—nearly five times the lifetime emissions of the average American car (and that includes manufacture of the car itself)."[63] Models with lower degrees of confidence had smaller carbon footprints. If we want to create and detect manufactured speech with the scale, speed, and certainty that platforms demand, we will harm the Earth. How do we want to allocate speech certainty, given its ecological impact?

Finally, probability matters to free speech because it goes to the heart of what it means to realize and govern ourselves. If, as Hacking argues, "statistics has consequences for the ways in which we conceive of others and think of our own possibilities and potentialities," then probability is a key logic of humanity.[64] If the chance that our words spread or that we hear others depends on probabilistic systems, then we have a vested interest in seeing probability as a political technology that either helps or hinders our abilities to think, associate, deviate, adapt, resist, or act. And when we limit probability to one type of concept, one particular operationalization or set of values, we limit our ability to imagine new social arrangements. If a group is inchoate, under attack, or actively marginalized—"enclave publics" who need special protections—should those groups be treated differently by probabilistic systems?[65] Perhaps they should not have to endure *any* of Facebook's or

Google's errors. Should people with a large number of followers or subscribers *automatically* have to statistically demonstrate that they are *less* likely to disseminate harmful speech, because they are closer to traditional broadcasters? (How many followers, what would the statistical test be, and who would make these decisions?) Should data journalists use different levels of statistical certainty depending on the public importance of a story and the potential impacts of their findings on particular policies and populations?[66] Different theories of the public demand different understandings of probability.

Conclusion

Online speech systems exist at a scale that makes human oversight practically impossible. Their scale forces them into "operating actuarially" and using probability to manage that scale.[67] Currently, probability is most often the instrument of those who have vested interests in maintaining large-scale surveillance economies. Scale makes money, and probability enables scale.

But we should not fear probability, reject it outright as a tool of oppression or control, or relegate it to the narrow domain of statisticians and technologists. Rather, we should interrogate speech systems *as probabilities* that beg new questions. Are these the meanings and applications of probability we need? Who has enough knowledge to confidently predict these systems' behaviors, and how can the private and proprietary nature of probabilistic knowledge be challenged? Who suffers from false positives and false negatives? Which interests and beliefs are embedded in the Bayesian classifiers that categorize online speech? [68] Which types of errors are known and tolerated, how is risk distributed, and who has the institutional standing or technological power to challenge thresholds, reject error rates, and renegotiate categories? Should platforms be held accountable not just for errors but error *rates* and confidence thresholds? Should we issue "probability taxes" on those who distribute risk in socially unacceptable ways?

To think anew about probability and the role it plays in speech governance, we need to think anew about the "permanently beta" culture that seems to *require* failure as a marker of success, encouraging innovators to set error thresholds too low, to move too fast and break too many things.[69]

But not all errors are alike, and they can compound with disastrous consequences. In her study of the Challenger explosion, sociologist Diane Vaughan describes how NASA made "routine decisions" that "normalized technical deviation" so successfully and invisibly that it was blind to the catastrophic failure their ritualized acceptance of increasingly compounded risks caused.[70] Where in platforms—in their construction, use, and governance —are there understandings of normalized deviance, compound risk, and catastrophic failure? Are platform errors always simply actuarial inevitabilities, or are there some that are so morally unacceptable they might be existential challenges to platforms, their technological designs, and their business models?

We are only beginning to understand how and why to regulate speech platforms. Probability should be a central part of this conversation.

© 2019, Mike Ananny.

1　Paul Horwitz, First Amendment Institutions (2013); Frederick Schauer, *The First Amendment as Ideology*, *in* Freeing the First Amendment: Critical Perspectives on Freedom of Expression (David S. Allen & Robert Jenson, eds., New York Univ. Press 1995); Laurence Tribe, The Invisible Constitution (2008); Mark Tushnet, Taking the Constitution Away From the Court (2000).

2　Horwitz, *supra* note 1.

3　*Id*. at 15.

4　Frederick Schauer, *Towards an Institutional First Amendment,* 89 Minn. L.J. 1256, 1265 (2005).

5　Horwitz, *supra* note 1, at 18–19.

6　Jack M. Balkin, *Free Speech is a Triangle,* 118(7) Colum. L. Rev. 2011 (2018).

7　Tarleton Gillespie, Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions that Shape Social Media (2018); Kate Klonick, *The New Governors: The People, Rules, and Processes Governing Online Speech,* 131 Harv. L. Rev. 1598 (2017); Jose van Dijck, Thomas Poell & Martjin de Waal, The Platform Society (2018).

8　Tarleton Gillespie, *Governance of and by Platforms, in* SAGE Handbook of Social Media 254–278 (Jean Burgess, Thomas Poell & Alice Marwick, eds., 2017).

9　Robert Gorwa, *What is Platform Governance?*, 22 Info., Commc'n. & Soc'y 854 (2019).

10　Kate Crawford & Tarleton Gillespie, *What is a Flag For? Social Media Reporting Tools and the Vocabulary of Complaint,* 18 New Media & Soc'y, 410 (2014); Stuart R. Geiger, *Bot-Based Collective Blocklists in Twitter: The Counterpublic Moderation of Harassment in a Networked Public Space,* 19 Info., Commc'n. & Soc'y 787 (2016); Sarah T. Roberts, Behind the Screen (2019).

11　Gorwa, *supra* note 9; David Kaye, A Human Rights Approach to Platform Content Regulation: Report of the Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression, A/HRC/38/35 (Apr. 6, 2018), https://freedex.org/a-human-rights-approach-to-platform-content-regulation/.

12　P.J. DiMaggio & W.W. Powell, The New Institutionalism in Organizational Analysis 14 (1991).

13　W.R. Scott, Institutions and Organizations: Ideas, Interests, and Identities (4th ed. 2013).

14　Donna Haraway, *Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective,* 14(3) Feminist Stud. 575 (1988).

15　P.J. DiMaggio, *Constructing An Organizational Field As a Profession Project: U.S. Art Museums, 1920–1940*, *in* The New Institutionalism and Organizational Analysis 267–292 (P.J. DiMaggio and W. Powell, eds., Univ. of Chicago Press, 1991).

16　Bourdieu and the Journalistic Field (Rodney Benson and Eric Neveu, eds., 2005).

17   Robyn Caplan & danah boyd, *Isomorphism Through Algorithms: Institutional Dependencies in the Case of Facebook,* 5 Big Data & Soc'y 1 (Feb. 14, 2018), https://journals.sagepub.com/doi/full/10.1177/2053951718757253; Philip M. Napoli, *Automated Media: An Institutional Theory Perspective on Algorithmic Media Production and Consumption,* 24(3) Commc'n. theory 340 (2014).

18   Mike Ananny, Networked Press Freedom: Creating Infrastructures For a Public Right to Hear (2018).

19   Geoffrey C. Bowker & Susan Leigh Star, Sorting Things Out: Classification and Its Consequences (1999); John Law, *Technology & Heterogenous Engineering: The Case of Portuguese Expansion, in* the Social Construction of Technological Systems 111–134 (Wiebe E. Bijker, Thomas P. Hughes & Trevor Pinch, eds., MIT Press, 1987); Sergio Sismondo, An Introduction to Science and Technology Studies (2nd ed. 2009); Lucy Suchman, Plans and Situated Actions: The Problem of Human-Machine Communication (1987).

20   This is a much abbreviated discussion of a rich area of scholarship. On the history of probability, the development of chance, and its application to a variety of sociotechnical domains, *see, e.g.* Louise Amoore, The Politics of Possibility: Risk and Security Beyond Probability (2013); Phaedra Daipha, Masters of Uncertainty: Weather Forecasters and The Quest for Ground Truth (2015); William Deringer, Calculated Values: Finance, Politics, and The Quantitative Age (2018); Gerd Gigerenzer et al., The Empire of Chance: How Probability Changed Science and Everyday Life (2013); Donald MacKenzie, Inventing Accuracy: An Historical Sociology of Nuclear Missile Guidance (1990); Theodore M. Porter, The Rise of Statistical Thinking: 1820–1900 (1986); Theodore M. Porter, Trust in Numbers (1995); Frederick Schauer, Profiles, Probabilities, and Stereotypes (2009); Stephen M. Stigler, The History of Statistics: The Measurement of Uncertainty before 1900 (1990).

21   Gigerenzer, *supra* note 20.

22   *Id*. at 7.

23   Ian Hacking, The Taming of Chance 2 (1990). On the politics of categories and conventions, *see* Bowker & Star, *supra* note 19; S.L. Star, *Power, Technologies, and the Phenomenology of Conventions: On Being Allergic to Onions, in* A Sociology of Monsters: Essays on Power, Technology, and Domination 26 (J. Law, ed., 1991); Marion Fourcade and Kieran Healy, *Classification Situations: Life-chances in the Neoliberal Era,* 38(8) Accounting, Org. and Soc'y 559; on the role of technologies in constructing populations, *see* John Law, *Seeing Like a Survey,* 3(2) Cultural Sociology 239; James C. Scott, Seeing Like a State (1999).

24   Gigerenzer, *supra* note 20, at 41.

25   Ian Hacking, *Making Up People* in Reconstructing Individualism 161–171 (T.L. Heller, M. Sosna & D.E. Wellbery, eds., 1985).

26   Gigerenzer, *supra* note 20.

27   *Id*. at 48.

28   *Id*.

29   *Id*. at 50–51.

30   Arunabh Ghosh, *Lies, Damned Lies, and (Bourgeois) Statistics: Ascertaining Social Fact in Midcentury China and the Soviet Union,* 33(1) Osiris 149, 155 (2018).

31   *Id*. at 165.

32   On this point, *see* Dan Bouk, How Our Days Became Numbered: Risk and The Rise of The Statistical Individual (2015); Sarah Igo, The Averaged American: Surveys, Citizens, and The Making of a Mass Public (2007).

33   Ian Hacking, An Introduction to Probability and Inductive Logic (2001); Jeffrey N. Rouder, Julia M. Haaf & Frederik Aust, *From Theories to Models to Predictions: A Bayesian Model Comparison Approach*, 85(1) Commc'n. Monographs 41 (2018).

34   For excellent discussions of the dangers of using "big data" in social media contexts, *see* Eszter Hargittai, *Is Bigger Always Better? Potential Biases of Big Data Derived from Social Network Sites*, 659(1) The Annals of Am. Acad. of Pol. & Soc. Sci. 63 (2015); Dhavan V. Shah, Joseph N. Cappella & W. Russell Neuman, *Big Data, Digital Media, and Computational Social Science: Possibilities and Perils*, 659(1) The Annals of Am. Acad. of Pol. & Soc. Sci. 6 (2015).

35   Hacking, *supra* note 23.

36   S. Gupta et al., *The Anatomy of a Large-Scale Experimentation Platform*, IEEE Int'l Conf. on Software Architecture (2018); Ya Xu et al., *From Infrastructure to Culture: A/B Testing Challenges in Large Scale Social Networks*, Proc. of the 21st ACM SIGKDD Int'l Conf. on Knowledge Discovery & Data Mining (2015).

37   Sarah Myers West, *Censored, Suspended, Shadowbanned: User Interpretations of Content Moderation on Social Media Platforms*, 11(20) New Media & Soc'y 4366 (2018).

38   Monika Bickert, *Defining the Boundaries of Free Speech on Social Media, in* The Free Speech Century 254–271, 269 (Lee C. Bollinger and Geoffrey R. Stone, eds., 2018). Twitter's Del Harvey makes a similar point about scale and error: "Given the context of the scale we're dealing with, if you're talking about a billion tweets, and everything goes perfectly right 99.999% of the time, then you're still talking about 10,000 tweets where everything might not have gone right." Tarleton Gillespie, *Platforms are Not Intermediaries*, 2 Geo. L. Tech. Rev. 198, n. 1 (2018).

39   Stefan Wojcik et al., *Bots in the Twittersphere*, Pew Res. Ctr. (Apr. 9. 2018), http://www.pewinternet.org/2018/04/09/bots-in-the-twittersphere/.

40   Madeline Lamo and M. Ryan Calo, *Regulating Bot Speech*, 66 UCLA L. Rev. 988 (2019). Sites like TweetBotOrNot and Botometer produce scores of how likely a Twitter account is a bot. (TweetBotOrNot reports that there is a 67 percent chance that my own account, @ananny, is a bot.)

41   Russell Brandom, *Deepfake Propaganda is Not a Real Problem*, The Verge (Mar. 5, 2019), https://www.theverge.com/2019/3/5/18251736/deepfake-propaganda-misinformation-troll-video-hoax.

42   Tessa Lyons, *Hard Questions: What's Facebook's Strategy for Stopping False News?*, Facebook (May 23, 2018), https://newsroom.fb.com/news/2018/05/hard-questions-false-news/.

43   *Continuing Our Work to Improve Recommendations on YouTube*, YouTube (Jan. 25, 2019), https://youtube.googleblog.com/2019/01/continuing-our-work-to-improve.html.

44   *Id.*

45   Will Knight, *An AI that Writes Convincing Prose Risks Mass-producing Fake News*, MIT Tech. Rev. (Feb. 14, 2019), https://www.technologyreview.com/s/612960/an-ai-tool-auto-generates-fake-news-bogus-tweets-and-plenty-of-gibberish/.

46   A growing body of literature shows the complexity of governing platform speech, tracing the various policy tensions and cultural dynamics of platform content moderation. It is an openly debated question whether platforms could impose such bans. Some suggest that bans are impossible because platforms are technologically unable to automatically identify suspect content, adding more moderators will make implementing standards even more inconsistent, and no combination of algorithmic and human moderation can be fast enough to meet the speed and scale of a nearly global set of users instantaneously distributing content. The debate suggests that platforms have created a scale of speech that they are simply unable to

govern reliably or that their business models prevent them from seriously and systematically curbing the designs they have created to encourage viral content distribution. *See, e.g.* Crawford and Gillespie, *supra* note 10; Gillespie, *supra* note 7; Gorwa, *supra* note 9; David Kaye, Speech Police: The Global Struggle to Govern the Internet (2019); Roberts, *supra* note 10; Nicolas P. Suzor et al., *What Do We Mean When We Talk About Transparency? Toward Meaningful Transparency in Commercial Content Moderation*, 13 Int'l J. of Commc'n. 1526.

47 Maggie Fick and Paresh Dave, *Facebook's Flood of Languages Leave it Struggling to Monitor Content,* Reuters (Apr. 22, 2019), https://www.reuters.com/article/us-facebook-languages-insight-idUSKCN1RZoDW.

48 Ken Walker, *Four Steps We're Taking Today to Fight Terrorism Online,* Google (June 18, 2017), https://blog.google/around-the-globe/google-europe/four-steps-were-taking-today-fight-online-terror/.

49 Jake Kanter, *Leaked Emails Reveal Facebook's Intense Internal Discussion over Alex Jones' "Anti-Semitic" Post on Instagram*, Insider (Mar. 28, 2019), https://www.insider.com/facebook-emails-reveal-discussion-about-alex-jones-instagram-account-2019-3.

50 *Id.*

51 Jane Guyer, *Percentages and Perchance, in* Legacies, Logics, Logistics: Essays in the Anthropology of the Platform Economy 140, 148 (2016).

52 Cade Metz and Mike Isaac, *Facebook's A.I. Whiz Now Faces the Task of Cleaning It Up. Sometimes That Brings Him to Tears,* N. Y. Times (Mar. 17, 2019), https://www.nytimes.com/2019/05/17/technology/facebook-ai-schroepfer.html.

53 Martin Kaste, *Facebook Increasingly Reliant on A.I To Predict Suicide Risk*, NPR (Nov. 17, 2018), https://www.npr.org/2018/11/17/668408122/facebook-increasingly-reliant-on-a-i-to-predict-suicide-risk.

54 Dan Muriello et al., *Under the Hood: Suicide Prevention Tools Powered by AI,* Facebook Code (Feb. 21, 2018), https://code.fb.com/ml-applications/under-the-hood-suicide-prevention-tools-powered-by-ai/.

55 Mason Marks, *Artificial Intelligence for Suicide Prediction,* Balkinization (Oct. 30, 2018), https://balkin.blogspot.com/2018/10/artificial-intelligence-for-suicide.html.

56 Catherine Card, *How Facebook AI Helps Suicide Prevention,* Facebook Newsroom (Sept. 10, 2018), https://newsroom.fb.com/news/2018/09/inside-feed-suicide-prevention-and-ai/.

57 Jacob Snow, *Amazon's Face Recognition Falsely Matched 28 Members of Congress With Mugshots,* ACLU (July 26, 2018), https://www.aclu.org/blog/privacy-technology/surveillance-technologies/amazons-face-recognition-falsely-matched-28.

58 Davey Alba, *Amazon Rekognition Falsely Matched 28 Members of Congress With Arrest Mugshots,* BuzzFeed News (July 26, 2018), https://www.buzzfeednews.com/article/daveyalba/amazon-rekognition-facial-recognition-congress-false.

59 *Amazon Rekognition Developer Guide: CompareFaces,* Amazon (2019), https://docs.aws.amazon.com/rekognition/latest/dg/API_CompareFaces.html.

60 T. Dryer, *Algorithms under the Reign of Probability,* 40 IEEE Annals of the Hist. of Computing 93.

61 Tim Wu, *Is The First Amendment Obsolete?,* Knight First Amend. Inst. (Sep. 2017), https://knightcolumbia.org/content/tim-wu-first-amendment-obsolete.

62 Emma Strubell et al., *Energy and Policy Considerations for Deep Learning in NLP,* 57th Annual Meeting of the Ass'n for Computational Linguistics (2019).

63   Karen Hao, *Training a Single AI Model Can Emit as Much Carbon as Five Cars in Their Lifetimes*, MIT Tech. Rev. (June 6, 2019), https://www.technologyreview.com/s/613630/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/.

64   Hacking, *supra* note 23, at 6.

65   Catherine R. Squires, *Rethinking the Black Public Sphere: An Alternative Vocabulary for Multiple Public Spheres,* 12 Commc'n. Theory, 446–468 (2002).

66   For a smart discussion of the role journalists in general, and data journalists in particular, play in constructing public notions of certainty, *see* C.W. Anderson, Apostles of Certainty: Data Journalism and the Politics of Doubt (2018).

67   Schauer, *supra* note 20, at 6.

68   Bernhard Rieder, *Scrutinizing an Algorithmic Technique: The Bayes Classifier as Interested Reading of Reality,* 20 Info., Commc'n. & Soc'y 100 (2016).

69   Gina Neff, Venture Labor (2012).

70   Diane Vaughan, The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA (1996).

MIKE ANANNY is an Associate Professor of Communication and Journalism at USC Annenberg.

FILED UNDER ESSAYS AND SCHOLARSHIP