BERKMAN KLEIN CENTER
FOR INTERNET & SOCIETY
AT HARVARD UNIVERSITY

Research Publication No. 2018-7
October, 2018

Common-Knowledge Attacks on Democracy

Henry Farrell
Bruce Schneier

23 Everett Street • Second Floor • Cambridge, Massachusetts 02138
+1 617.495.7547 • +1 617.495.7641 (fax) • http://cyber.law.harvard.edu/ •
cyber@law.harvard.edu

Common-Knowledge Attacks on Democracy

by

Henry Farrell
Professor of Political Science and International Affairs at George Washington University

Bruce Schneier
Fellow, Berkman Klein Center for Internet and Society
Fellow and Lecturer, Belfer Center and Digital HKS, Harvard Kennedy School

October 2018

Abstract

Existing approaches to cybersecurity emphasize either international state-to-state logics (such as deterrence theory) or the integrity of individual information systems. Neither provides a good understanding of new "soft cyber" attacks that involve the manipulation of expectations and common understandings. We argue that scaling up computer security arguments to the level of the state, so that the entire polity is treated as an information system with associated attack surfaces and threat models, provides the best immediate way to understand these attacks and how to mitigate them. We demonstrate systematic differences between how autocracies and democracies work as information systems, because they rely on different mixes of common and contested political knowledge. Stable autocracies will have common knowledge over who is in charge and their associated ideological or policy goals, but will generate contested knowledge over who the various political actors in society are, and how they might form coalitions and gain public support, so as to make it more difficult for coalitions to displace the regime. Stable democracies will have contested knowledge over who is in charge, but common knowledge over who the political actors are, and how they may form coalitions and gain public support. These differences are associated with notably different attack surfaces and threat models. Specifically, democracies are vulnerable to measures that "flood" public debate and disrupt shared decentralized understandings of actors and coalitions, in ways that autocracies are not.

Introduction

In 2014, presumed Russian hackers sought to compromise key aspects of Ukraine's elections. Notably, the targets included the systems used to communicate the election results to newspapers. As a Ukrainian official described the attack: "Offenders were trying by means of previously installed software to fake election results in the given region, and in such a way to discredit general election results of elections of the President of Ukraine."[1]

In 2016, the Internet Research Agency, a company based in St. Petersburg, began to post false content on US social media that seemed intended to stir up controversy, division, and disagreement on the facts among its readers, to the point of trying to create both protests and counter-protests over the same issues.[2] Many scholars doubt whether these attacks had large-scale consequences for behavior,[3] but they plausibly worsened a general sense of paranoia, doubt, and confusion among people who were increasingly unsure what their fellow citizens believed, and (as the debate over Internet manipulation began) which of them were fellow citizens, and which foreign trolls or automated processes.[4]

Both these attacks are attacks on *common political knowledge*: the consensus beliefs that hold political systems such as democracies together.[5] Election security does not simply involve physical infrastructure, such as ballots and polling booths. It also involves roughly consensual expectations about how the system works, who won and who lost, and so on. If an attacker does not penetrate the physical election infrastructure, but does successfully subvert the shared expectations around the election, she can nevertheless succeed.[6]

To work properly, democracies require this kind of broad agreement across many questions. Democracies delegate core aspects of decision making to ordinary citizens, and to politicians and parties who struggle for citizens' support. Politi-

1 Shackelford, Scott J., Schneier, Bruce, Sulmeyer, Michael, Boustead, Anne, Buchanan, Ben, Deckard, Amanda N. C. et al. (2017). "Making Democracy Harder to Hack: Should Elections Be Classified as 'Critical Infrastructure'?" University of Michigan Journal of Law Reform, 50(3), 629–668.

2 See Shane, Scott, and Mazetti, Mark (2018). "The Plot to Subvert an Election: Unraveling the Russia Story So Far." *New York Times.*

3 Benkler, Yochai, Faris, Rob, and Roberts, Hal (2018). *Network Propaganda: Manipulation, Disinformation and Radicalization in American Politics.* New York: Oxford University Press.

4 Chen, Adrian (2016). "The Real Paranoia-Inducing Purpose of Russian Hacks." *New Yorker.*

5 We note for readers familiar with game theory that our understanding of "common knowledge" is less demanding than the formal definition they are familiar with. Everybody does not need to know what everyone else knows, and so on. Instead, for us, common knowledge is the roughly shared set of social beliefs about how the system works, who the actors are, and so on, which helps to order politics. We stress its coordinating role in this paper: other, more sociologically inclined, accounts might stress legitimacy instead. We leave for later debate the extent to which these differing accounts might better or worse capture actual political dynamics.

6 As Bruce Schneier wrote regarding election security: "Elections serve two purposes. The first, and obvious, purpose is to accurately choose the winner. But the second is equally important: to convince the loser." Schneier, Bruce (2018). "American Elections Are Too Easy to Hack. We Must Take Action Now." *Guardian.* See also Shackelford et al., "Making Democracy Harder to Hack."

cians and the citizens they represent will disagree over many topics. However, if the decentralized system of democracy is not to break down into chaos, then citizens and their representatives have to roughly agree about what they disagree about.[7] They have to be able to recognize who the different factions in society are and what their broad purposes are, and to believe that their political opponents will not seek to permanently dominate or destroy them, but instead will be subject to the same democratic limits as they are. Attacks that undermine these collectively held expectations will make it far harder for groups and parties to make coalitions, forge compromises, or engage in the rest of the grind of democratic politics.

Common-knowledge attacks can have critical consequences,[8] yet they are a poor fit with conventional national security approaches to cybersecurity. National security officials traditionally think about cybersecurity using Cold War concepts that were developed to understand nuclear weapons. They use ideas such as the offense–defense balance, conventional deterrence theory, and deterrence by denial.[9] They focus on the threats posed by nation-state adversaries. They consider how best to mitigate these threats in a low-information environment, both by manipulating information about capabilities and intentions, and — where appropriate — making credible threats against adversarial states.

Technologists, in contrast, start from a very different set of security assumptions. Broadly speaking, they are agnostic about whether the threats come from states or other actors. Instead, they focus on defending specific information systems: modeling potential threats that different kinds of actors might pose to these systems based on their characteristics. They want to understand the *attack surface* that attackers might exploit, and close off or mitigate the most serious vulnerabilities or the ones that widen the attack surface.[10] They try to design secure and reliable systems based on a deep understanding of how attacks and attackers operate.

Neither of these approaches is innately well-suited for analyzing common-knowledge attacks. On one hand, national security officials have a hard time using the traditional concepts of national security theory to analyze the threats exemplified by these attacks. Some aspects, such as the hacking of electoral databases or systems containing sensitive political information, fit traditional notions of state-on-state espionage and covert action. But the ways in which this hacked information has been used to stir up political controversy are a much poorer fit, and efforts to "flood"[11] social media with irrelevant and distracting content in order to compromise democratic debate do not fit at all.

7 We deliberately choose not to discuss here the thorny question of how much each individual citizen has to know for democracy to function properly, and how much of the work can instead be delegated to broader structures, such as political parties.

8 They include the attacks by Russia against the 2016 UK Brexit vote, the 2016 US election, and the 2017 French presidential election.

9 See for example Nye, Joseph S. (2011). "Nuclear Lessons for Cyber Security?" *Strategic Studies Quarterly*, 54(4), 18–38, "Deterrence and Dissuasion in Cyberspace." *International Security*, *41*(3), 44–71, Lindsay, Jon R. (2013). "Stuxnet and the Limits of Cyber Warfare." *Security Studies*, *22*(3), 365–404.

10 See for example Schneier, Bruce (2001). *Secrets and Lies: Digital Security in a Networked World.* Hoboken, NJ: John Wiley, Bellovin, Steven (2015). *Thinking Security: Stopping Next Year's Hackers.* Boston: Addison-Wesley Professional, Shostack, John (2014). *Threat Modeling: Designing for Security.* Hoboken, NJ: John Wiley.

11 On flooding, see Roberts, Margaret (2018). *Censored: Distraction and Diversion Inside China's Great Firewall.* Princeton, NJ: Princeton University Press.

Some national security analysts and scholars use concepts such as information warfare, or — more pejoratively — propaganda. This captures some aspects of these new forms of attack, but does less well at capturing others. These attacks tend to be more aimed at degrading than persuading; that is, at making democratic debate more difficult rather than attempting to change people's minds in a particular direction. While national security scholars have sought to analyze influence and "chaos" attacks as aspects of a common phenomenon, it is not clear that they fit well together. [12] Some of this writing is unduly alarmist: for example, Clint Watts warns of the threat of "Advanced Persistent Manipulators," claiming that "hacking people's computers...feels like child's play compared to the hacking of people's minds that has occurred on social media platforms the past four years."[13] Finally, it is hard to see how standard approaches to deterrence can provide a plausible solution to these attacks,[14] especially when they are not carried out by nation-state adversaries.

On the other side, technologists' understanding of these attacks is equally flawed. Attacks on election systems or on political parties' private servers are broadly similar to other attacks on information systems and can partly be mitigated by better threat modeling, better-designed systems, and more extensive security technologies. Some of the specific technological support structures of American democracy (most notoriously, voting machines) are highly vulnerable and in sore need of redesign according to commonly understood security principles.[15] However, technologists don't usually think systematically about broader knowledge systems and expectations. Instead, they focus on narrowly defined traditional information systems, such as servers and individual networks, and have little to say about the consequences of attacks for the broader fabric of democratic societies. While technologists can note the possibility of such effects, they do not have any good means of evaluating the associated risks.

One way to remedy this gap is to extend the logic of national security further, so that it looks to explain and counter a variety of nontraditional and nonmilitary threats that have consequences for freedom, liberty, and democracy. Realist scholars such as Jack Goldsmith have long been skeptical about the US effort to extend its liberal and democratic values via the Internet, believing that this radically underestimates the differences between different nation-states and the capacity of those states to defend their interests against outside incursions.[16] More recently, on his own and in collaboration with Stuart Russell, Gold-

12 Lin, Herb, and Kerr, Jaclyn (2018). "On Cyber-Enabled Information/Influence Warfare and Manipulation." In Paul Cornish (Ed.), *Oxford Handbook of Cyber Security*. New York: Oxford University Press, Paul, Christopher, and Matthews, Miriam (2016). *The Russian "Firehose of Falsehood" Propaganda Model: Why It Might Work and Options to Counter It*. Santa Monica, CA: Rand Corporation.

13 P. 7, Watts, Clint (2018). "Advanced Persistent Manipulators and Social Media Nationalism: National Security in a World of Audiences" (Aegis Series Papers 1812). Palo Alto, CA: Hoover Institution.

14 See Goldsmith, Jack (2016). "The DNC Hack and the (Lack of) Deterrence." *Lawfare*.

15 See Blaze, Matt (2017). *Testimony*. Proceedings from US House of Representatives Committee on Oversight and Government Reform, Subcommittee on Information Technology and Subcommittee on Intergovernmental Affairs Hearing on Cybersecurity of Voting Machines, Washington DC.

16 See in particular Goldsmith, Jack, and Wu, Tim (2006). *Who Controls the Internet: Illusions of a Borderless World*. New York: Oxford University Press.

smith has argued that attacks against democracy demonstrate the profound flaws of the Internet freedom agenda. Specifically, he claims that the US "pro freedom" bias against censorship and regulating the commercial actors who dominate the Internet has created major national security vulnerabilities, and rendered the US incapable of responding to profound new threats.[17]

Goldsmith's skepticism about the dominance of commercial actors is well founded. However, the national security perspective is systematically blinkered in ways that make it hard to assess the appropriate means to defend democratic practices against incursions. When viewed from the perspective of national security, most forms of freedom — almost by definition — are also potential vulnerabilities. This means that the national security approach has enormous difficulties in assessing the appropriate trade-offs that are needed to guarantee a well-functioning democracy. Intelligent versions of the national security perspective, such as Goldsmith's, at least note the need for these trade-offs and the difficulty in striking them. Cruder versions may end up identifying the freedoms that they are purportedly supposed to defend as windows of vulnerability that need to be closed.

In this paper, we argue that extending the technical approach to cybersecurity provides a different — and we believe more useful — way of understanding the problem. Like national security thinkers, we begin from the level of the nation-state. But like technologists, our analysis focuses on the informational aspects of the nation-state. Specifically, we are interested in the different ways that democracies and autocracies organize themselves, and the kinds of coordination that they need to function effectively. The technical approach does not currently address questions of collective knowledge, but there is no reason in principle why it cannot.

Hence, we scale up the technologists' approach to cybersecurity, so that rather than thinking about specific information systems within democracy, it approaches democracies and autocracies as *information systems,* and then asks questions such as what is their respective attack surface, which likely threat models they face, and how do they (or can they) seek to mitigate risks? The technical approach to cybersecurity is precisely intended to strike trade-offs between ensuring that the information system is usable and accessible, and minimizing and mitigating the inevitable vulnerabilities that go together with usability and accessibility.

We sketch out a simple framework that both identifies key differences between autocratic and democratic systems, and provides a roadmap for future research and policy measures.

As far as we know, no one has previously undertaken this kind of analysis. In one sense, that is surprising, given its plausible relevance and value. In another, it is not surprising at all. There is existing research literature on the informational trade-offs or "dictators' dilemmas" that

17 See Goldsmith, Jack (2018). *The Failure of Internet Freedom.* Miami FL: Knight Foundation, Goldsmith, Jack, and Russell, Stuart (2018). "Strengths Become Vulnerabilities: How a Digital World Disadvantages the United States in Its International Relations" (Hoover Institution Aegis Papers 1806). Palo Alto, CA: Hoover Institution.

autocrats face, in seeking to balance between their own need for useful information and economic growth, and the risk that others can use available information to undermine their rule.[18] There is no corresponding literature on the informational trade-offs that democracies face between desiderata like availability and stability.

Common and Contested Political Knowledge

All societies face important trade-offs between two kinds of political information: (a) *common political knowledge* — the knowledge that everyone in the political system needs to share in order for it to function, and (b) *contested political knowledge* — the knowledge that is contestable, where people may disagree.

Common political knowledge involves a body of information that people in many societies broadly, if loosely, agree on. This is the roughly shared knowledge that allows for decentralized political coordination. For example, in a stable autocracy, people agree on who the rulers are and what their legitimating ideology involves. In a stable democracy, citizens agree that their votes count, and that election results reflect the actual distribution of opinion in society — even if only roughly and imperfectly.[19]

In a democracy, the common political knowledge about institutional rules and the range of other actors does not have to be encyclopedic, but it does have to provide a sufficient shared understanding of how politics works to provide general social stability. One of the crucial insights of an academic body of work associated especially with Barry Weingast and his co-authors is that "open access orders," like those of the advanced industrial democracies, require a variety of stabilizing informal expectations if they are to work in a coordinated way.[20] Equally, autocracies rely on a broad set of shared expectations to function well. As Russell Hardin notes, no government is strong enough to impose its will on its population if the population decides not to cooperate with it.[21] Common political knowledge is what provides stabilizing expectations in both instances.

18 See Roberts, *Censored*, pp. 23–25, for discussion of the literature, Kalathil, Shanthi, and Boas, Taylor (2003). *Open Networks, Closed Regimes: The Impact of the Internet on Authoritarian Rule.* New York: Carnegie Endowment for an early articulation of the logic of the dictator's dilemma, and Hollyer, James, Rosendorff, Peter, and Vreeland, James (2018). *Information, Democracy, and Autocracy: Economic Transparency and Political (In)Stability.* New York: Cambridge University Press for a useful recent treatment.

19 On democratic stability and expectations, see Fearon, James (2011). Self-Enforcing Democracy. *Quarterly Journal of Economics*, *126*(4), 1661-1708. There is of course a well-established literature in social choice on the impossibility of reaching outcomes that truly reflect people's individual preferences, given a set of reasonable criteria. There are further inevitable distortions that arise from institutions, political parties, and so on. Nonetheless, democracies roughly reflect people's different beliefs and perspectives in ways that autocracies do not.

20 North, Douglass, Wallis, John Joseph, and Weingast, Barry R. (2009). *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History.* Princeton, NJ: Princeton University Press, Hadfield, Gillian, and Weingast, Barry R. "What is Law? A Coordination Model of the Characteristics of Legal Order." *Journal of Legal Analysis*, 4(2), 471–514, Carugati, Federiga, Hadfield, Gillian, and Weingast, Barry R. (2015). Building Legal Order in Ancient Athens. *Journal of Legal Analysis*, 7(2), 291–324.

21 Hardin, Russell (1990). "The Social Evolution of Cooperation." In Karen Cook and Margaret Levi (Eds.), *The Limits of Rationality*. Chicago: University of Chicago Press.

This loose agreement on what everyone "knows" coexists with a quite different, and even contrary, form of knowledge: the information dispersed in the political disagreements within a given society, or, as we call it, contested political knowledge.[22] This is the political knowledge that emerges from the tensions between the different goals and perspectives of various actors and groups in society. For example, people in a democracy may disagree on questions such as the role that government should play in the economy, or whether there should be tariffs or free trade, or how the government should conduct its foreign policy.

All societies have real or potential political factions and actors, or coalitions of actors, each with its own specific goals. Very often, these goals conflict with each other; for one actor or coalition to achieve its goal is to frustrate another's. These differing goals are commonly associated with different cognitive styles of problem solving, and different beliefs about what the most important problems are. Politics, then, is the process through which these group conflicts over goals and problem-solving styles, and rankings of problems are expressed, mediated, and suppressed.

This distinction is poorly understood by academics, let alone policy makers, because existing work tends to focus on one or the other and not the relationship between them. Thus, for example, some strategic accounts of politics focus on the need to generate common expectations that allow for broad social coordination even in decentralized societies.[23] Others instead emphasize the degree of diverse knowledge and beliefs within society, and the problems and/or benefits that arise therefrom.[24]

It is obvious that society organized around a government cannot survive without common political knowledge. What is less obvious is that contested political knowledge is also valuable. Just as, for scholars of biological evolution, the level of information in a species is contained in its genetic diversity, the extent of reasonable political disagreement in a society is a rough index of the information that society possesses. Complex social problems are best solved when multiple, diverse perspectives can be applied to them, each perspective potentially disclosing an aspect of the problem that is invisible to others.[25]

22 Here we develop ideas articulated first in Farrell, Henry, and Shalizi, Cosma R. (2015). "Pursuing Cognitive Democracy." In Danielle Allen and Jennifer Light (Eds.), *From Voice to Influence: Understanding Citizenship in a Digital Age* (pp. 211–231). Chicago: University of Chicago Press (we are grateful to Cosma for basic insights that inform our broader arguments).

23 See in particular Hadfield and Weingast, "What is Law?" While Hadfield and Weingast acknowledge the importance of "idiosyncratic" understandings as a spur to creative economic interactions, their theory insulates these understandings from the self-enforcing institutions that they see as fundamental to the stability of open access orders.

24 Knight, Jack, and Johnson, James (2011). *The Priority of Democracy: Political Consequences of Pragmatism.* Princeton, NJ: Princeton University Press, Levy, Jacob (2018). *Justice In Babylon.* Unpublished paper.

25 Page, Scott. (2007). The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies. Princeton, NJ: Princeton University Press, Lazer, David, and Friedman, Allan (2005), "The Parable of the Hare and the Tortoise: Small Worlds, Diversity, and System Performance." Kennedy School of Government Working Paper No. RWP05-058.

Democracies and Autocracies as Information Systems

The relationship between common political knowledge and contested political knowledge differs sharply across autocracies and democracies. Many important aspects of politics have incompatible needs for common and contested political knowledge. Where democracies require that a certain aspect of political knowledge be contested, autocracies may require that it be commonly held, and vice versa.

Democracies draw upon the disagreements within their population to solve problems.[26] Partisanship is the set of political disputes between collective actors that have organized around different interests and, typically, different associated ways of understanding the collective problems faced by a given society. In a well-functioning democracy, each such group vies for political influence by persuading voters that its way of understanding problems and associated solutions is the best one. This is to the democracy's benefit.[27] It provides a mechanism through which a polity can harness the diversity of perspectives within it, the better to solve complex problems.

This requires contestation over who the rulers should be, and what broad social goals they will seek to implement. Political parties and other collective actors hope that they (or their allies) will be in control for a given period, and each vies against others to win public support to that end. This also requires long-term uncertainty over who will be in charge and able to set policy goals. In successful democracies, the rules of democratic competition provide the uncertainty that drives parties to creatively reconfigure problems and propose solutions so as to appeal to voters and perhaps win future elections.

Two different kinds of common political knowledge among these collective actors and the general public are key to the proper functioning of democracy.

The first involves political institutions — for example, rules over elections, succession of power, and so on — that channel conflict in a democratic society. Institutions are only effective when the relevant people in a given institution agree on what they are, how they work, and what their consequences are.[28] If political actors are to have the incentive to compete and, even more importantly, to concede when they have lost, they need to have common knowledge that the voting system is fair, and that a short-term loss may still allow them to compete and win in the future.[29]

26 Farrell and Shalizi, "Pursuing Cognitive Democracy."

27 Rosenblum, Nancy L. *On the Side of the Angels: An Appreciation of Parties and Partisanship,* Princeton University Press, 2010, Farrell and Shalizi, "Pursuing Cognitive Democracy."

28 See Jack Knight, *Institutions and Social Conflict* (Cambridge University Press 1992), although note that institutions invariably will involve a rough rather than a complete consensus on what the rules mean. See further Danielle Allen, Henry Farrell and Cosma Shalizi, *An Evolutionary Account of Institutional Change.* Unpublished paper.

29 See Przeworski, Adam (2018). *Why Bother with Elections*. Hoboken, NJ: John Wiley for a good recent overview of elections and alternation of government.

They will also need to share common knowledge that their domestic adversaries are broadly committed to the democratic process, so that they need not fear indefinite domination or worse when they and their allies lose.

The second involves common knowledge over the range of actors, beliefs, and opinions in the society. If new interests and new parties are to come into being, compete, and either fail or flourish, they will need to have a reasonable understanding of who the other political actors are, what their interests are, and where they clash with or converge with their own. While much attention is paid to the generic costs of collective action, the shared knowledge that allows political actors to identify and coordinate with potential allies, attract voters, and so on is just as important, and perhaps more so.

In successful democratic societies, knowledge is decentralized across the wide variety of collective actors whose consent and willingness to constrain their activities is necessary for the system to work.[30] This is essential, since ordinary citizens play a significant role in political decision making instead of just handing authority to a central power elite.[31]

Autocracies adopt a very different approach to common and contested knowledge. In contrast to democracies, they require common political knowledge about who is in charge, and what their social goals are, as a basic condition of stability. There may be internal contestation between different factions within the elite, but such contestation is often clandestine, and is carefully insulated from the public realm, so as not to destabilize the shared expectations that anchor regime stability.

This explains the great lengths that autocracies often go to manipulate shared expectations, and to support useful public beliefs. Autocracies benefit — as democracies do not — from what political scientists have described as "pluralistic ignorance" or "preference falsification," under which people only have private knowledge of their own political beliefs and wants, without any good sense of the beliefs and wants of others.[32]

For example, Marc Lynch has noted that the Tunisian autocracy was one of the "most heavily censored states on earth."[33] It relied on an information environment in which public displays of support for the regime were mandated, informing on friends, neighbors and family was common, and dissidents were tortured and punished, so

30 For relevant models, see James Fearon, "Self-Enforcing Democracy," Little, Andrew, Tucker, Joshua, and LaGatta, Tom (2015). Elections, Protest, and Alternation of Power. *Journal of Politics*, *77*(4), 1142–1156, Hollyer, Rosendorff, and Vreeland, *Information, Democracy, and Autocracy.* In all of these models, the possibility of mass protest plays a key role in providing rulers with sufficient incentive to relinquish office.

31 Rosenblum, *On the Side of the Angels*, stresses the democratic problems associated with defining who the "people" are in exclusionary ways so as to preempt future changes in a democracy's self-conception.

32 Kuran, Timur (1997). *Private Truths, Public Lies: The Social Consequences of Preference Falsification.* Cambridge, MA: Harvard University Press, King, Gary, Pan, Jennifer, and Roberts, Margaret (2017). "How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument." *American Political Science Review*, 111(3), 484–501.

33 P. 305, Lynch, Marc (2011). "After Egypt: The Limits and Promise of Online Challenges to the Authoritarian Arab State." *Perspectives on Politics*, 9(2), 301–310.

that it was difficult for individuals to know how truly unpopular the regime had become: all they could see was their own private unhappiness, and the public support shown by others.[34] Even if an autocratic government is broadly detested, it may remain in power so long as the public does not realize how broadly detested it is.

Autocracies do not require common political knowledge about the efficacy and fairness of elections. In Valerie Bunce's pungent distinction, while democratic countries provide certainty about the political process and uncertainty over outcomes, authoritarian countries provide uncertainty about process and certainty over outcomes.[35] Many authoritarian regimes conduct elections, both as a legitimating sop and to provide themselves with some information as to the distribution of views within their population. However, they typically show no compunction in manipulating the results to ensure that the regime and its supporters triumph. The actual workings of electoral institutions — and representative institutions more generally — are likely to be opaque to ordinary citizens and outsiders. Authoritarian regimes also often ensure that the "rules of the game" of politics are hidden, or open to manipulation or revision, in order to ensure that upstarts can't use those rules to organize against them.

Autocratic regimes will typically benefit from contested political knowledge about nongovernmental groups and actors in society. Again, efficacious long-term collective action does not merely rest on simple technologies that make it cheaper for people to organize. It also requires a relatively sophisticated understanding of the variety of different political actors both inside and outside the ruling coalition, and the distributed support among the population for these actors and their differing agendas. However, self-organizing collective actors are more likely to be a challenge than a resource to autocratic regimes, since such actors may become powerful enough to form coalitions that challenge the regime, leading to a transition either to another form of rule or a new autocracy with different actors in charge.

Rather than allowing common political knowledge regarding the preferences of the population and the variety of political actors to be shared among actors in a decentralized order, such regimes will try to maintain monopolistic control. This forestalls new collective interests from organizing, and makes it harder for existing interests to coalesce into a challenging coalition. To preserve the stability of their own rule, they will look to prevent independent interests from having sufficient appeal to broad segments of the population, and to prevent the population from being attracted to — and associating themselves with — independent interests.

Hence, they will act to limit common political knowledge about potential groupings in the society, their likely levels of support, and the possible coalitions they can form among each other. Thus, for example, the extensive Chinese social media censorship system is less focused on shaping the expression of public opinion (which may be valu-

34 International Crisis Group (2011). *Popular Protest in North Africa and the Middle East (IV): Tunisia's Way* (106). Brussels, Belgium: International Crisis Group.

35 We are grateful to Adam Segal for informing us of this distinction, and Valerie Bunce for confirming it.

able to the state under some circumstances) than on preventing citizens and others from organizing around particular causes.[36]

To be sure, autocracies may want accurate information *for themselves* about political beliefs within the population, so as to keep track of their legitimacy and ensure their long-term stability. Thus, authoritarian regimes such as the former Soviet Union, even while they kept tight control of public information through extensive censorship and surveillance, kept track of public beliefs through extensive survey polling, the results of which were only available to elite party leaders.[37] Modern autocracies such as China similarly rely on public opinion surveys, as well as on social media as an index of broad public sentiment.[38]

Thus, there are crucial differences between democracies' and autocracies' respective approaches to information and knowledge. These differences mean that forms of information that may be stabilizing for one may be destabilizing for the other.

The Attack Surfaces of Autocracies and Democracies

The differences described have important consequences for security. Authoritarian regimes are potentially vulnerable to information attacks that challenge their monopoly on common political knowledge, either by undermining preference falsification or by disseminating knowledge in ways that allow other collective actors to organize and form coalitions to challenge the regime. They also are vulnerable to attack vectors that turn contested knowledge and uncertainty among potential regime adversaries about their levels of popular support, ability to form coalitions, and so on, into usable common political knowledge.

Democratic regimes, in contrast, are vulnerable to information attacks that widen contested political knowledge so that it spills over into disagreements over the common political knowledge that democracy needs to operate. They are similarly vulnerable to attack vectors that turn contested knowledge over who will rule and to what ends, into common political knowledge that permanently advantages a specific faction and associated set of social goals.[39] Finally, they are vulnerable to attacks on the common political knowledge shared by groups, factions, and parties about their respective goals, levels of political support, and potential coalitions, as well as to attacks on shared expectations about the fairness of the political system. Because this knowledge is decentralized, it is easier to destabilize through certain kinds of attacks. The level and kind of vulnerabilities will

36 King, Pan, and Roberts, "How the Chinese Government Fabricates Social Media Posts."

37 Dimitrov, Martin (2014). "Tracking Public Opinion Under Authoritarianism: The Case of the Soviet Union During the Brezhnev Era." *Russian History*, 41(3), 329–353.

38 Dimitrov, Martin (2015). Internal Government Assessments of the Quality of Governance in China. *Studies in Comparative International Development*, 50(1), 50–72.

39 We note in passing that this provides one possible way of understanding the internal (and perhaps in some cases, to some limited degree, external) attacks on democratic knowledge and expectations that have helped turn countries such as Hungary and Poland into populist democracies.

differ across different democratic regimes. Most notably, where common political knowledge is already frail, it will be easier for adversaries to engineer further attacks.

This difference helps us understand how policy measures that increase the stability of one form of regime may decrease the stability of another. The history of the last two decades has demonstrated how open information flows that benefited democratic regimes were viewed by authoritarian regimes as an existential threat, because they might transform regime-supporting contested political knowledge into regime-threatening common political knowledge. Only recently have we started to understand how the same information flows that benefit autocracies can be weaponized against democracies, turning regime-supporting common political knowledge into regime-undermining contested political knowledge.

Until quite recently, Western academics and policy makers shared a broad consensus about the destabilizing consequences of open information flows for autocratic regimes. This consensus dated from the mid-1990s, when libertarians such as John Perry Barlow claimed that the Internet would undermine tyrannical rule and extend freedom.[40] It also extended to the left and the non-libertarian right. In different decades, both Bill Clinton and Hillary Clinton have made similar claims.[41] The George W. Bush administration instituted a program of spending millions of dollars to provide technological assistance to anti-censorship activists, which was continued under the Obama administration.[42]

A proper understanding of the importance of common political knowledge and contested political knowledge helps explain both (a) why open information flows were regarded as an uncomplicatedly good thing by most Western observers, and (b) how they could have specific negative consequences for authoritarian regimes. These flows seemed to support the decentralized common political knowledge of democratic regimes rather than undermining it, providing better information to both political groups and voters about the broad contours of democratic politics, the range of actors, and the public support that they had. Internet communications technologies further provided the means for new groups to identify their shared interests and self-organize, helping the Howard Dean campaign, the left-leaning Netroots, Tea Party Republicans, and Black Lives Matter to circumvent traditional institutional barriers.[43]

40 Barlow, John Perry (1996). *Declaration of the Independence of Cyberspace.* San Francisco: Electronic Frontier Foundation (republished).

41 Clinton, William Jefferson (2000). "Remarks at the Paul H. Nitze School of Advanced International Studies." *Johns Hopkins SAIS*, Clinton, Hillary Rodham (2012). Internet Freedom and Human Rights. *Issues in Science and Technology, 28*(3), 45–52.

42 Kiggins, Ryan D. (2015). "Open for Expansion: US Policy and the Purpose for the Internet in the Post-Cold War Era." *International Studies Perspectives*, 16(1), 86–105, Goldman, "Strengths Become Vulnerabilities."

43 Johnson, Stephen B. (2008). "Two Ways to Emerge, and How to Tell the Difference Between Them." In Jon Lebkowsky and Mitch Ratcliffe (Eds.), *Extreme Democracy.* extremedemocracy.com, Farrell, Henry (2006). "Bloggers and Parties: Can the Netroots Reshape American Democracy?" *Boston Review*, Carney, Nikita (2016). "All Lives Matter, but So Does Race: Black Lives Matter and the Evolving Role of Social Media." *Humanity and Society*, 40(2), 180–199. There is some disagreement about the extent to which the Tea Party was an organic grassroots movement, but see Skocpol, Theda, and Williamson, Vanessa (2012). *The Tea Party and the Remaking of American Conservatism.* New York: Oxford University Press on the reliance of local activists on online tools such as MeetUp.

In contrast, such flows had potential destabilizing consequences for authoritarian regimes. The preference falsification that regimes such as Tunisia relied upon could be undone by social media like Facebook, which was not then censored or widely monitored. As the common political knowledge about the regime's stability started to unravel, it became easier for individuals to come together and challenge it in public.

This was reinforced in Tunisia and elsewhere by the creation of new forms of common political knowledge where previously there had been contested political knowledge. As new technologies substantially lowered the costs of collective action, it became easier (in principle) for people to organize in groupings outside state structures.[44] As these groups became more aware of other groups, and their various goals and levels of public support, they could begin to form coalitions, which in time could challenge and even potentially topple the regime. In many cases, it turned out that these coalitions did not lead to a democratic transition, but the prospect of long-term failure provided little comfort to threatened authoritarian leaders. The enthusiasm of democratic leaders for technology-fueled challenges to authoritarianism helped fuel paranoia among leaders who saw themselves as targeted, so that Vladimir Putin, for example, described the Internet as a "CIA project."[45] There is reason to believe that Russia's hacking attacks during the US elections were in part motivated by the desire for retaliation.[46]

Contrary to these hopes and fears, the Internet and communications technologies have no inherent bias toward freedom.[47] Indeed, authoritarian regimes proved adept at quickly turning new technologies to their purposes. On one hand, they started to use social media as an alternative means of safely gathering information about public preferences.[48] On the other, they learned how to shut down and drown out potentially dissident voices.[49] Many authoritarian regimes began to supplement fear-based forms of censorship with "friction" aimed at dissuading ordinary members of the public from looking for certain kinds of information through increasing the costs, and "flooding" public forums so as to disrupt decentralized public knowledge building and coalition building.[50] They furthermore sought increasingly to exclude foreign NGOs focused on open society- and democracy-related issues from their domestic politics.

44 Shirky, Clay (2008). Here Comes Everybody: The Power of Organizing without Organizations. New York: Penguin.

45 See McAskill, Ewan (2014). "Putin Calls Internet a 'CIA Project' Renewing Fears of Web Breakup." *Guardian*.

46 Ioffe, Julia (2018). "What Putin Really Wants." *Atlantic*.

47 Morozov, Evgeny (2011). *The Net Delusion: The Dark Side of Internet Freedom*. New York: Public Affairs, Farrell, Henry (2012). "The Consequences of the Internet for Politics." *Annual Review of Political Science*, 15(1), 35–52.

48 Gunitsky, Seva (2015). "Corrupting the Cyber-Commons: Social Media as a Tool of Autocratic Stability." *Perspectives on Politics*, *13*(1), 42–54.

49 Tucker, Joshua A., Theocharis, Yannis, Roberts, Margaret, and Barberá, Paolo. (2017). "From Liberation to Turmoil: Social Media and Democracy." *Journal of Democracy*, 28(4), 46–59.

50 Roberts, *Censored*.

Adrian Chen describes the results in Russia:

> ...after speaking with Russian journalists and opposition members, I quickly learned that pro-government trolling operations were not very effective at pushing a specific pro-Kremlin message — say, that the murdered opposition leader Boris Nemtsov was actually killed by his allies, in order to garner sympathy. The trolls were too obvious, too nasty, and too coordinated to maintain the illusion that these were everyday Russians. Everyone knew that the Web was crawling with trolls, and comment threads would often devolve into troll and counter-troll debates. The real effect, the Russian activists told me, was not to brainwash readers but to overwhelm social media with a flood of fake content, seeding doubt and paranoia.[51]

The systemic consequences of such measures inside Russia were to make the formation of common political knowledge impossible outside the parameters set by the government. If US libertarians claim that the best antidote to bad speech is more speech, Putin's government discovered that the best antidote to more speech was even more bad speech. Thus, authoritarian governments such as China, and semi-authoritarian regimes such as Russia moved quickly to mitigate vulnerabilities in their information systems, through disrupting the ability of both domestic and international actors to turn regime-favoring contested political knowledge into regime-undermining common political knowledge about the genuine state of public beliefs.

Tools such as flooding can stabilize authoritarian and semi-authoritarian regimes, but are likely to disrupt the common knowledge that is necessary to the successful functioning of democracy.

This explains the Russian influence attacks against the US in 2016. There is substantial reason to believe that the information ecology of US democracy had already been substantially weakened by internal forces.[52] Specifically, a right-wing media ecology had evolved that was separate from and antagonistic to the mainstream, which rapidly conveyed extreme arguments from the fringes of the system to the center, serving as a force-amplifier of lies.

These media structures plausibly created wide vulnerabilities. However, the information that has emerged via the Mueller indictment of individuals associated with the Internet Research Agency (IRA) and other sources sketches out an account of flooding attacks that fits closely with our arguments.[53]

First, some of the attacks focused directly on undermining belief in the electoral system. As the Mueller indictment describes the attacker's intentions: "By in or around May 2014, the ORGANIZATION's strategy included interfering with the 2016 U.S. Presidential Election," with the stated goal of "spread[ing] distrust towards the candidates and the political system in general."

51 Chen, "The Real Paranoia-Inducing Purpose."

52 Benkler, Faris and Roberts, *Network Propaganda*.

53 United States v. Internet Research Agency et al., available at https://www.justice.gov/file/1035477/download.

This likely explains why Russian actors helped propagate rumors that Hillary Clinton was guilty of vote fraud as well as probing the vulnerability of online US electoral records. Their probable intentions were not to fix the vote but to create enough paranoia over the possibility that the vote had been fixed that Hillary Clinton's legitimacy would have been seriously damaged, had she been elected as president. "Guccifer 2.0," a pseudonymous identity used by Russian intelligence, claimed just before the election that "the Democrats may rig the elections on November 8. This may be possible because of the software installed in the FEC networks by the large IT companies. As I've already said, their software is of poor quality, with many holes and vulnerabilities."[54] Plausibly, the attackers did not expect Trump to be elected president. Instead, they wanted a United States that was sufficiently divided against itself that a President Hillary Clinton would have difficulty in governing, let alone taking decisive actions abroad.

Second, other attacks are aimed more generally at creating division between different groups, damaging and breaking up existing coalitions, and preventing new ones from forming. For example, on the Affordable Care Act:

> The Russian effort moved easily between supporting and opposing the health law depending on the political moment. Pro-ACA tweets peaked around the spring of 2016, possibly aimed at fostering division between Mrs. Clinton and her presidential primary rival, Sen. Bernie Sanders (I., Vt.). Anti-ACA tweets intensified in mid-2017 as Republicans mounted their push to repeal the law, apparently seeking to capitalize on the emotions generated by that effort. "Let Obamacare crash & burn. Do not bail out insurance companies," said a tweet from an IRA-linked account called JUSMASXTRT on Aug. 28, 2017.[55]

Such attacks disrupt democracy by degrading citizens' and groups' shared political knowledge about allied and adversarial groups within society, fomenting confusion about the goals of those groups, and the level and kind of support that they enjoy. By increasing the levels of noise, flooding attacks degrade the decentralized common political knowledge that provides people with a rough overall map of politics, and make it more difficult to organize around collective interests or to build coalitions across interest groups. People may also come to believe that fringe beliefs are more widespread in the population than in fact they are, widening the political debate

54 https://guccifer2.wordpress.com/2016/11/04/info-from-inside-the-fec-the-democrats-may-rig-the-elections/.

55 Armour, Stephen, and Overberg, Paul (2018). "Nearly 600 Russia-Linked Accounts Tweeted about the Health Law." *Wall Street Journal.*

so that it includes perspectives that enjoy little actual public support. Finally, they may substantially increase paranoia, which further degrades knowledge and makes political action harder. If people believe that they are surrounded by trolls and bots, they are more likely to be distrustful of others (especially others with different beliefs) and less likely to engage in dialogue or effective political action.[56]

Democracy Defenses

A better understanding of the informational requirements of democracy does not merely help us to understand the attack surface better. It also provides a clearer understanding of how to bolster security in democracies.[57] Specifically, it implies a series of broad priorities. We sketch them out in this section in order to spur discussion, which may lead in time to a properly developed policy agenda.

The first among these, and likely the least controversial, is to better defend the common political knowledge that democracies require to function.[58] We do not understand enough about the institutions that help support this knowledge, and cannot yet provide a detailed list of defenses. At this point, more research is required (we hope that this paper will help spur such research). Very obviously, voting systems are a crucial source of political knowledge. Not only successful compromise of such systems, but also attacks aimed at weakening public beliefs and expectations surrounding the fairness of these systems, can seriously damage political common knowledge. Thus, it is important to supplement the valuable recommendations of the National Academies of Science for improving the security of the voting system itself with a more specific understanding of the public perceptions and beliefs surrounding voting, so as to frustrate more subtle attacks on expectations.[59]

However, there are other important sources of common knowledge that present less obvious security risks. For example, the US Census was instituted precisely to provide the public and experts with a common understanding of the demographics of the US population, so as to better aid apportionment of political power and public policy. Attacks that are aimed at weakening the Census — or public expectations surrounding it — may also damage the informational supports of democracy, by excluding portions of the population, by reshaping beliefs about the relative role and influence of different demographic groups in society, and so on. Mapping out other such

56 These effects involve indirect consequences. The direct effects of disinformation on people's beliefs and behavior may be relatively weak. Barberá, P. et al. (2018). *Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature.* Palo Alto, CA: Hewlett Foundation.

57 Our arguments also have implications for understanding autocratic politics, including perhaps developing a better understanding of how different autocracies will be affected by knowledge attacks. We note this as a suggestion for future research and debate.

58 See also Nye, Joseph S. (2018). "How Sharp Power Threatens Soft Power: The Right and Wrong Ways to Respond to Authoritarian Influence." *Foreign Affairs.*

59 National Academies of Science. (2018). *Securing the Vote: Protecting American Democracy.* Washington DC: National Academies of Science Press.

institutions, and an agenda of specific measures to protect them, presents an urgent challenge for researchers and policy makers.

The second set of priorities involves the nexus between inside and outside groups. The extensive literature in computer security describes how attacks by outsiders can be greatly facilitated when insiders provide specific information and (sometimes inadvertent) help. While cooperation is generally a good thing, there are also problematic forms of cooperation — and institutional changes that render such problematic forms easier to get away with.

Thus, for example, recent legal changes and changes in interpretation of the law make it far easier for foreign actors to work together with domestic actors in clandestine ways. This may plausibly damage democracy, such as by funding campaigns that are aimed directly at spreading public disinformation. Concerns have been expressed about possible Russian funding for Marine Le Pen's National Front party in the French general elections in 2017, the "Leave" campaign in the United Kingdom's Brexit referendum, and other movements that seem likely to contribute to general political instability in rich democracies.[60] If these concerns are valid, they highlight specific vulnerabilities that may be widened further by current difficulties in tracking political spending on social media. The general move away from publicly disclosed political funding and toward nontransparent forms of political spending creates obvious vulnerabilities. In future, flooding attacks may combine cross-border and internal campaigns to much greater effect. One obvious implication of our framework is that dark money structures are not only ethically problematic, but create obvious security vulnerabilities, which could be mitigated through far stricter reporting requirements.

Similar arguments apply to large-scale social media companies such as Facebook and YouTube (owned by Google). These companies' business models make it easier to conduct clandestine information operations with little external visibility. As scholars like Zeynep Tufekci have argued, they also may exacerbate the damage of common-knowledge attacks, such as through algorithms that maximize on user "engagement" and hence drive users toward material that reinforces conspiratorial thinking.[61] The plausible responses to such problems entail major reform, whether by far greater regulation, the transformation of these companies into public utilities, or their being broken up. Each of these options presents a different set of benefits and drawbacks.

These arguments should not be applied indiscriminately. There are a wide variety of non-problematic and democratically beneficial relations between those who are citizens (insiders) of a given democratic system and those who are outside. As global interdependence increases, creating new problems that span borders, some forms of cross-border cooperation are not only helpful to

60 See Gatehouse, Gabriel (2017). "Marine Le Pen: Who's Funding France's Far Right?" *BBC News* and Reuters (2018). "Brexit-backer Arron Banks Denies Fresh Allegations of Russia Links."

61 Tufekci, Zeynep (2018). "YouTube. The Great Radicalizer." *New York Times.*

democracy, but positively essential for it — global warming being the most obvious example.[62]

Finally, and potentially most controversially, the computer security literature makes no strong distinction between insider and outsider effects, except insofar as they have different opportunities to compromise the information system. Again, our framework of analysis has broad implications for how to defend democracy, suggesting that institutions that allow insiders to compromise common democratic knowledge can heavily damage democracy. Notably, however, there is a tension between the need to maintain common democratic knowledge of the kinds that we have described, and the need to allow the contested democratic knowledge that is necessary for successful democratic problem solving. We do not even pretend to offer a complete account of how this tension should best be managed. Instead, we point out that certain features of the US political system (e.g., the widely observed disparities of political influence between rich and poor)[63] both hamper the contestation and political debate that we argue is necessary to democratic success, and plausibly enhance the risk from insider threats. While reforming these features is an enormously ambitious political agenda, it does not present obvious trade-offs between security and democratic functioning.[64]

Conclusions

In this paper, we make three basic claims. First, we argue that we currently do not have a good theory of the kinds of influence attacks that have afflicted the US and other democracies over the last few years. Both national security and technical security approaches to cybersecurity have notable deficiencies in understanding how these attacks operate. Second, we argue that substantially — and even radically — expanding the technical security approach provides the best and most appropriate means to developing such a theory. If we treat national political regimes as information systems, we can better understand their attack surfaces and threat models. Third, we use these combinations to explain the different attack surfaces of autocracies and democracies, demonstrating, for example, how measures that improve stability in autocracies may have destabilizing consequences in democracies, and vice versa. We believe that this account better captures the potential policy trade-offs in defending against such attacks than the most plausible alternative — developing and applying the national security perspective.

This last requires more justification. The most comprehensive national security account of influence attacks that we are aware of is Jack

62 Farrell, Henry, and Knight, Jack (2018). *John Dewey's Lessons for Interdependence.* Unpublished Paper.

63 Bartels, Larry (2016). *Unequal Democracy: The Political Economy of the New Gilded Age.* Princeton, NJ: Princeton University Press.

64 Concerns about these disparities has typically been the focus of the left. However, a new body of work builds on classical liberalism to a broadly similar set of conclusions. See in particular Lindsey, Brink, and Teles, Steven (2017). *The Captured Economy: How the Powerful Enrich Themselves, Slow Down Growth and Increase Inequality.* New York: Oxford University Press.

Goldsmith and Stuart Russell's recent essay, "Strengths Become Vulnerabilities." As the authors describe their argument:

> Our central claim is that the United States is disadvantaged in the face of these soft cyber operations due to constitutive and widely admired features of American society, including the nation's commitment to free speech, privacy and the rule of law; its relatively unregulated markets; and its deep digital sophistication. These strengths of American society create asymmetric vulnerabilities in the digital age that foreign adversaries, especially in authoritarian states, are increasingly exploiting. ...We do not claim that the disadvantages of digitalization for the United States in its international relations outweigh the advantages. But we do present some reasons for pessimism about the United States' predicament in the face of adversary cyber operations.

We do not contend that Goldsmith and Russell's pessimism is completely unwarranted. Defending democracy against these kinds of attacks will be a Herculean labor.[65] However, we think that Goldsmith and Russell's pessimism is exaggerated by the difficulty that the national security perspective has in thinking systematically about the appropriate trade-offs. When the perspective of national security is extended to influence operations (or, as we prefer, common-knowledge attacks), just about every opening looks like a vulnerability.

Goldsmith and Russell do not conclude that this means that all those vulnerabilities need to be closed. Instead, they propose that we are faced with a series of unpleasant trade-offs between what makes American society admirable, and what is necessary to protect it from outside encroachment. However, they have no useful metric to determine how difficult trade-offs ought be struck. This is in part, we suspect, because the national security approach is not designed for people to think systematically about the internal benefits of openness. Indeed, standard realist accounts assume that what happens within states and what happens between them are analytically entirely separate.

Here, the computer security approach provides a better foundation. Since information systems need to be open to input if they are to be useful, computer security analysts are trained to think systematically about the trade-offs between openness and security and then balance the requisite equities. First, one needs to understand what a given information system is supposed to do. Then, one needs to weigh the forms of input and access that are necessary for functioning against the attack vulnerabilities that different modes of input and access provide. Typically, one cannot provide comprehensive solutions, but — through design and experiment — one can mitigate the vulnerabilities associated with openness to the point that the benefits outweigh the risks. First, one looks to pluck the low-hanging fruit, by closing vulnerabilities that have few or no benefits. Then, one carefully assesses the benefits and drawbacks of the more complex trade-offs between openness and vulnerability.

65 See also Joseph Nye, "How Sharp Power Threatens Soft Power."

This is why an informational account of democracy is so important to mitigation. Without it, one cannot understand how democracy is supposed to operate and, hence, one cannot assess the trade-offs. The informational understanding that we present here emphasizes the way that democracy can draw on diverse sources of information. This means that democracies can potentially do better than autocracies over the long run, to the extent that they are better able to use the disagreements and diverse information they contain to solve complex collective problems. However, this also confronts them with the serious challenge of ensuring that the common political knowledge that provides stability is not overwhelmed by internal disagreements. Attacks that seek to widen internal disagreement so that it implicates common political knowledge can have very serious consequences.

Obviously, many of our policy priorities flow from this understanding of democracy. A different understanding might be the foundation of different prescriptions. However, we note that our understanding is sufficiently broad to be shared by an emerging set of arguments on the center-right of American debate as well as the left. And even those who disagree sharply with our premises and conclusions may draw some benefit from using a similar approach to analysis to think systematically about the informational foundations of democracy and its relationship to security.

Acknowledgments

We are grateful to Ross Anderson, Yochai Benkler, Sheri Berman, Maria Farrell, Rob Faris, Martha Finnemore, Art Goldhammer, Anna Grzymala-Busse, Alex Grigsby, Herb Lin, Joseph Nye, Molly Roberts, and Adam Segal for their helpful comments on earlier versions of this paper. Some of this material was presented at a meeting of the Twenty-First Century Trust in Strasbourg, France, September 14–17, 2018.

A permalink to this report is available here:

https://cyber.harvard.edu/story/2018-10/common-knowledge-attacks-democracy