Democracy’s Dilemma

(this is a work in progress. Please don’t cite without checking with us first)

Henry Farrell and Bruce Schneier

Until recently, American policy makers, pundits, and scholars believed that the Internet would undermine authoritarian rulers. Bill Clinton famously likened China’s attempts to control the Internet was “sort of like trying to nail Jell-O to the wall.” George W. Bush and Barack Obama mostly agreed. The US government bet hundreds of millions of dollars on the proposition that the Internet would help spread democracy and build a better world. Perhaps the influx of new ideas would reduce the government’s strangleholds on political discourse. Perhaps the free flow of information would help oppressed citizens realize how much they all hated their government. Or perhaps it would help nourish the beginnings of an independent civil society, or, more simply, just make it cheaper and easier to organize protests.

Today we live in more despairing times. It isn’t just that authoritarians are using these technologies to bolster their rule. Pipe-dreams have morphed into night-terrors. Instead of undermining dictatorship, the Internet is undermining democracy. Donald Trump’s campaign team claim that they were able to use targeted advertising to win the presidency. Russians and others are attempting to manipulate elections across the world. Social media is not only riddled with bots and propagandists, but warped by the perception that bots and propagandists have replaced real disagreements with artful manipulation. A new consensus is emerging: that American democracy is less a political system than a free fire zone in a broader information war, a wilderness of mirrors where nothing is real and everything is suspect.

There are real reasons to worry. Democratic competition can spill over into destabilizing internal conflict and manipulation more easily than we realized. Yet this is not enough reason to give up hope. The Arab Spring wasn’t the twilight of dictatorship, and today isn’t necessarily the twilight of democracy. Democracy may have systematic informational weaknesses that we didn’t know about before, and both domestic and foreign actors are taking advantage of those weaknesses. Still, there is a great deal of ruin in our political system: even if things are going badly, they are not as catastrophic as they sometimes appear in the news. It is likely possible to mitigate these problems, or even fix some of them.

We first need to figure out what they are. That’s harder than it might seem. In an ideal world, social scientists, political theorists, philosophers, and information security experts would have well-developed theories that they could use to separate justified from fanciful fears. We’re not in that world. Instead, everyone -- academics and experts included -- are scrambling to figure things out on the fly. Most of our broad intellectual frameworks have assumed that democracy’s relationship to information is unproblematic, or at least that democracy is far better able to deal with uncontrolled information flows than other kinds of government. We need new frameworks to understand a world in which that’s not necessarily true.

****

One valuable source of ideas is traditional research on authoritarian politics. Scholars and policy thinkers have spent a lot of time thinking about what is called the Dictator's Dilemma – the tradeoff autocrats face between political stability and open information flows. Openly available information allows autocratic governments, including autocratic ones, run better: by alerting government officials to what's really happening among their citizens, by allowing markets to function properly, and by identifying low level corrupt officials that stymie policy and make citizens unhappy. A government without accurate information about its country and people risks unnecessarily enacting and enforcing unpopular, inefficient, and ineffective policies.

However, freely available information can also undermine a despot. It makes it easier for regime opponents to organize protests and appeal to citizens. It allows citizens to figure out how deeply unpopular a detested regime is, building up their confidence in shared collective action. A government that allows information to flow freely risks losing power.

Authoritarian regimes constantly need to balance this trade-off, and generally choose to protect their own power at the expense of efficiency. This helps explain why so many authoritarian regimes have seen the open Internet as a fundamental attack on their stability: by opening up political information to actors in civil society, it makes it possible for those actors to undermine or even topple the regime. It also explains why they usually have not tried to get rid of the Internet entirely, because they fear hurting economic growth and damaging the ability of the higher levels of government to keep track of what the lower levels are doing. The result is that authoritarians have adopted a variety of mitigating strategies, which are intended to preserve as much as possible of the Internet's benefits for growth and efficiency, while blocking many of its politically dangerous features.

What we need now is to understand a corresponding problem: Democracy's Dilemma. Democracies benefit from the free flow of information in more fundamental ways than autocracies. Information not only allows markets to function, and helps government work more efficiently; it is essential to the decentralized decision making by ordinary citizens that democracy relies on. But these same information flows can be manipulated to undermine democracy by allowing the spread of propaganda and pseudo-facts -- all made more efficient by the Internet and automation. The dilemma that democracy faces is that the open forms of input and exchange that it relies on can be weaponized to inject falsehood and misinformation that erode democratic debate.

Questions abound. Which aspects of free information flows enhance the workings of democracy, allowing it to take advantage of the diverse knowledge of citizens, solve problems, and provide political stability across groups who might otherwise harm each other? Which aspects, instead, can be weaponized to undermine democracy? And how can we best identify mitigating strategies, which would preserve as many as possible of democracy's benefits, while preventing chaos and harmful instability?

Figuring this out will require a lot of hard work. We need to bridge very different ways of thinking. Social scientists and political theorists understand some aspects of democracy quite well, including questions of democratic legitimacy: what makes people believe in democracy and accept democratic verdicts, even when they are not what they want in the short term. For the most part, they haven't thought systematically about the flows of information and knowledge that democracies require to function properly. We can learn lessons from our scholarship surrounding the Dictator's Dilemma – but we need to be careful. We cannot crudely assume that democracies and autocracies face identical, or even similar kinds of informational threats.

Another useful source of scholarship is information security. Information security scholars are really good at modeling information systems and ways in which they can be sneakily attacked. They often think in terms of mitigation rather than perfect solutions: weighing for example, users' legitimate need for open access to information against the risks that attackers may exploit it. This leads to a more pragmatic understanding of security as a series of acceptable tradeoffs, rather than an all-or-nothing war against the enemy. They don't however, have a very good sense of how information systems interact with people's beliefs, which means that they have difficulty in understanding large scale attacks that look to leverage general social and political vulnerabilities.

Democracy's Dilemma can be described in the language of political theory and social science. What are the fundamental informational underpinnings of democracy, and how can they be undermined through manipulation? What do citizens need to know and believe in common for democracy to be sustainable, and how can that be manipulated and undermined? Democracy's Dilemma can also be understood in the language of computer security. Under what circumstances can information attacks destabilize democracies? Who are the most likely and most potent attackers? Or more succinctly, what is the attack surface of democracy, and what is the most plausible threat model?

To truly understand Democracy's Dilemma, we need to bring these two threads together. It's unlikely that Facebook memes, or Russian disinformation attacks are effective in spreading fake beliefs in the crude ways that the media talks about. The Russian GRU almost certainly did not win the election for Donald Trump. The American public is less bitterly divided than Twitter suggests. The problems are more subtle. Democracies probably require certain kinds of shared information if they are to work well. This may possibly be disrupted by deliberate attacks, by the unintended side-effects of political strategies, or by the internal workings of new forms of social media.

To understand if and how democracy is being disrupted – and to create mitigating policies – we need insights from computer security specialists, who have thought deeply about vulnerabilities in information systems. We also need insights from political theorists, political scientists, sociologists, legal scholars and others, who have thought about the ways in which social beliefs and expectations help to stabilize democracies. And somehow, we need to bring them together into a shared framework.

Such a framework would help us better understand the extent of the threat and the appropriate solutions. There is a danger, visible in much of the public debate, that the threat of information attacks on democracies will lead policy makers to enact watered-down versions of the same kinds of censorship and information control that authoritarian countries use to maintain their own political stability. Not only would that be harmful to civil liberties, it would fundamentally misunderstand the problem. Democracies have fundamentally different informational vulnerabilities than those of autocracies.

****

In a famous book among political scientists, Adam Przeworski argues that "democracy is a system of ruled open-endedness, or organized uncertainty" (p.13), where actors compete in elections over who should rule. One key question is why losers comply with electoral outcomes, rather than just deciding to subvert the electoral institutions, taking to the streets, or mounting a coup. Przeworski argues that they comply because they believe that even if they lose today, they realize that they have more to gain in the long term. As long as they have a chance of winning in the future, losing under democracy is be better than the non-democratic alternative.

Przeworski's insight is twofold. First, that the right kinds of uncertainty can have stabilizing democratic implications, and second that they will do so when harnessed by appropriate institutions, such as competitive elections. If parties could look into the future and know for sure that they would keep on losing, they might decide that their best strategy would be to overturn democracy.

One can could go further than Przeworski, as scholars like Jack Knight and James Johnson, and Henry Farrell and Cosma Shalizi have argued. Electoral politics can harness the deep disagreements between people with radically different understandings of the world and turn it into an engine of creativity, as different parties, representing different understandings of the problems that society faces and the right ways to solve it vie peacefully for the control of government.

Nonetheless, Przeworski's crucial insight travels. Uncertainty over democratic outcomes (who rules) and predictability over the democratic processes through which these outcomes are reached can reinforce each other in beneficial ways. Furthermore, it highlights a crucial difference between democracies and autocracies. As Valerie Bunce has observed, democracies provide certainty about the process and uncertainty about the outcomes, while autocracies provide uncertainty about the process and certainty about the outcomes. In democracies, people know how the government is elected, but they don't know who will form the government, while in autocracies they know who the government is, but the process through which the government comes into being is manipulated and opaque.

Scholars have been interested in the conditions under which democracy can be a self-enforcing equilibrium. They have asked how authoritarian regimes can be disrupted by new kinds of information. They haven't often asked how information might disrupt the workings of democracy, since democracy, vigorous disagreement and open information flows have seemed to go hand-in-hand. Yet implicit in Przeworski's and Bunce's arguments is the possibility that some kinds of disagreement can damage democracy. If autocracies can get into trouble when people start thinking that the political outcome (the regime's indefinite survival) is in doubt, democracies can go wrong when people become uncertain about the process.

We used these ideas to motivate an initial set of arguments, where we argue that democracies and autocracies are vulnerable to different kinds of destabilizing attacks. Rather than invoking some broad notion of uncertainty versus certainty, we talk to differences in what people agree on, and what they disagree about. This leads us to emphasize the tension between what we call *common political knowledge* and *contested political knowledge.*

In our argument, common political knowledge means the kinds of loosely shared knowledge that all the members of a society agree on. Contested political knowledge means the kinds of knowledge over which people disagree. Our crucial claim is that autocracies and democracies need antithetical kinds of common political knowledge and contested political knowledge to thrive.

This may seem like an abstract claim, but it has very practical implications. What it means is that autocracies will want to have common knowledge over who is in charge and their associated ideological or policy goals. Everyone should agree who the rulers are, what they want, and that they (or their chosen assigns) will remain in power for a very long time. However, contested knowledge can be valuable for autocracies too. Autocracies will be more stable if they can generate contested knowledge in the general population over who the various political actors in society are, and how they might form coalitions and gain public support. This kind of disagreement makes it more difficult for alternative coalitions to displace the regime.

In contrast, stable democracies will have contested knowledge over who is going to be in charge, and what they will hence do with their control over the government. As Przeworski suggests, this uncertainty about who may win future electoral contests helps both politicians and their supporters accept electoral defeat, giving them the incentive to keep on competing. However, democracies also require some kinds of diffuse common knowledge to do well. Specifically, there should be diffuse knowledge over who the major political actors are, how they may form coalitions and how they may influence policy and win elections by gaining sufficient public support. Barry Weingast and others have pointed out that open political orders like democracies decentralize much of the political action to citizens and to groups that emerge from civil society to contest politics. This also implies that citizens need at least a rough shared sense of politics, and what the rules of the game are, if the system is to work well.

These arguments help us understand the Dictator's Dilemma better. Increased information flows destabilize autocracies by making it more difficult for them to control what their subjects know and hindering them from organizing in ways that constrain or endanger the regime. In short, they replace contested political knowledge that stabilizes the regime with common political knowledge that undermines it. They help us to understand Democracy's Dilemma, too. The key implication is that information flows destabilize democracy when they replace stabilizing common political knowledge over process with destabilizing contested political knowledge.

This means, for example, that information flows can destabilize democracy if they succeed in undermining confidence in elections and other forms of preference aggregation. Such flows can erode the shared beliefs that make democracy more stable. Democracies have also turned to wider forms of public consultation to guide the policy making process; these too are potentially vulnerable to disruption. More subtly, these flows can also undermine the informational advantages of democracy – which stem from its ability to use elections and other means of information gathering to better guide policy.

Information flows that turn ordinary partisan rivalries into profound distrust, can be destabilizing, but under more limited conditions. For example, some people may believe that Donald Trump has been personally compromised by Russian agents, and is deliberately trying to undermine American democracy. Others may believe that socialists, atheists and Muslims are plotting to illegally remove Donald Trump from office. Such clashing beliefs can have serious consequences for people's willingness to accept that they have lost an election. After all, who would trust that a Russian catspaw, or that a coalition of radical America-haters will give up power willingly once it is gained. Even if there are not many people with such strong beliefs, second order beliefs about beliefs may destabilize democracy. I may think that the people on my side of politics are mostly reasonable, and worry that the people on the other side of the political spectrum might be as crazy as they seem. This may not only make me fear that they will cheat to win – it may make me want my own side to cheat, so that they don't get the chance.

We can understand these problems in terms of theories of democracy. The political science literature implies, even if it does not develop, arguments about the circumstances under which democracies may be vulnerable to certain kinds of information flows, precisely because their stability depends on some shared assumptions. We can also think about them in terms of information security. The most obvious vulnerability in the attack surface of democracy is the common political knowledge over process that allows democracies to decentralize key aspects of decision making to ordinary citizens.

There are two immediately plausible threat models. First, international actors may weaponize information flows to target the domestic politics of democracies that they see as strategic rivals or threats. They can even turn to tools that they have developed to shore up their own domestic security by creating confusion and disarray among citizens to do this, since the kinds of contested political knowledge that autocracies thrive on are likely to destabilize democracies. Second, domestic political actors may also disrupt collective political knowledge over process, either deliberately (where they believe that they will benefit from less democratic politics) or as a side-product of short term goals (where, for example, they disingenuously contest electoral results, with possible longer term implications for their supporters' trust in the electoral process).

****

The above is only a beginning, but we think that it is a valuable one. If we start thinking systematically about Democracy's Dilemma, we can begin to cut the problem down to size. First, we can begin to understand what the real threats are. Second, we can identify how they do not simply stem from foreign influence attacks, but from domestic actors too. Finally, we can start to think clearly about how best to respond to them.

If we are to take Democracy's Dilemma seriously, we can't just focus on Russian manipulation efforts. These efforts are real, and were more plausibly aimed at undermining confidence in an election that Trump was expected to lose than at helping Trump to win. Yet if they are important, they are only important as specific evidence of a more general set of informational vulnerabilities, where domestic actors are likely to cause bigger and more immediate problems.

Take, for example, three apparently unconnected recent controversies: efforts by groups supporting Doug Jones in the recent Alabama senatorial election to weaponize social media against conservatives, problems in the Federal Communications Commission (FCC) commenting process, and disagreements over the US Census. None of these involve the kinds of direct foreign attacks on US democracy that many commentators focus on. However, they all demonstrably illustrate different potential attack vectors against the common political knowledge that helps to stabilize democracy.

While there is disagreement over circumstances, extent, aims and lines of authority, it appears that groups supporting Doug Jones, the Democratic candidate in a heated 2016 Alabama senatorial election, used manipulative techniques on Facebook to target conservatives. Specifically, they created a "false flag" Facebook page aimed at dividing Republicans, for example by amplifying reports that the Republican candidate had pursued teenage girls, and suggesting that he was supported on social media by Russian bots. These actions illustrate an important dynamic of escalation. The effort appears to have been spurred in part by the desire to "fight fire with fire": the belief that the Russians had weaponized social media to help Trump get elected led some Democrats to experiment with similar tactics. This will likely lead Republicans in turn both to challenge the legitimacy of the 2016 election in 2020 ((What do you mean here? Republicans won in 2016. Why would they challenge it four years later?)), and to consider accelerating their own use of such techniques themselves in turn. The likely consequence of this spiral is increased disagreement and contestation over the legitimacy of elections as a by-product, rather than direct aim, of the actors involved.

In contrast, it is possible that efforts to hack the FCC's comment process in the recent fight over net neutrality were intended simultaneously to skew the process and to undermine public confidence in it. Net neutrality is universally disliked by large telecommunications companies, since it restrains them from using their privileged position to extract rents. However, it is highly popular with a mobilized and technically literate population of Internet users. When the FCC, under new chairman Ajit Pai, proposed to abandon net neutrality, it had to accept public comments. The result was a flood of legitimate comments from identifiable citizens, who were overwhelmingly in favor of keeping net neutrality, and a far larger flood of automatically generated comments, millions of which had erroneous or stolen email addresses, supporting Pai's plans to abandon it. This generated an overall vote that apparently favored getting rid of net neutrality, but only because of a deliberate generalized attack on the FCC's public commenting system. Even if the vote was obviously wrong, it quite effectively undermined the legitimacy of the commenting process in general, robbing the supporters of net neutrality of what would otherwise have been an overwhelming demonstration of apparent public support. Ongoing investigations by the New York Attorney General's office and the FBI have targeted telecommunications trade groups, lobbyists and advocacy organizations with subpoenas.

Finally, the US Census process is becoming a new information battleground. The Census plays a key role both in determining the distribution of seats in the House of Representatives and in the allocation of spending in a wide variety of government programs. It is an example of common political knowledge; everyone agrees on its accuracy That agreement is being undermined. The current administration has sought to include a new question about citizenship status, which would likely substantially depress the willingness of non-citizens, especially those with complex visa situations, to answer the census, leading to a likely substantial undercount of non-citizens. Others are circulating conspiracy theories on social media, trying to convince non-citizens and minorities not to fill out the census. Here, political actors are apparently deliberately seeking to undermine an important source of shared political knowledge in pursuit of particular political goals. These actions, combined with other problems of funding and mismanagement, may result in many believing that the Census is inaccurate, further resulting in beliefs that the government is "rigged" against them.

These three examples illustrate different vulnerabilities of democracy to specific informational effects and techniques. Attacks can undermine confidence in elections, disrupt channels of public commentary and input, and skew the institutions that we rely on to aggregate politically necessary knowledge. All of these disrupt common political knowledge, allowing uncertainty to destabilize democracy rather than stabilizing it.

Addressing these vulnerabilities requires a very different approach than the information war perspective that dominates public argument, which emphasizes deterrence, counter-attack and active defense. It also means avoiding any easy equivalence between the informational problems of autocracies and the informational problems of democracies. It is cheaper for autocracies to resort to censorship and information control because they rely less on decentralized choice and distributed public disagreement as engines of innovation.

We need to focus on protecting common political knowledge. Democracy's Dilemma does not involve a general breakdown of politics, but a specific set of threats to the core institutions that generate and disseminate the process knowledge that allows democracy to function. The best solution to this dilemma is to start finding ways to protect these institutions, while preserving their openness. This is the kind of design problem – striking tradeoffs between usability and vulnerability - that information security specialists have been thinking about for a very long time. However, practically addressing it will involve attention to phenomena such as legitimacy and social beliefs that information security specialists don't think about very often. Here, we need the knowledge of political scientists, political theorists, other social scientists and legal scholars. On the one hand, this will stretch the capacities of social science academics, who are unused to thinking about system design. On the other, it will require information security scholars to think about knowledge and belief systems that work in counter-intuitive ways.

Two cautions are appropriate. First, even with radical reforms, it will be impossible to eliminate uncertainty and disagreement about how institutions work. Institutions involve the application of abstract principles to generally messy real life situations, which always creates ambiguities that can be argued over. The purpose of protection should not be to eliminate disagreement about how institutions should be implied but to delimit it, so that it does not spill out to overwhelm democratic knowledge more generally. Second and related, shoring up common political knowledge about the fairness of the electoral system should obviously not lead to suppression of arguments about specific ways in which it is unfair, or about appropriate remedies to fix this unfairness.

Just the opposite is true. The general lesson is that the best way to tackle democracy's dilemma is to build appropriate and justified confidence that democratic processes indeed work as they ought. While this implies better and more visible public communication, it also implies institutions and systems that work more or less as they are supposed to. The claim that the best way to shore up democracy's vulnerabilities is to strengthen democratic institutions should seem obvious, but it's surprisingly uncommon in contemporary debate. A strong democracy is more likely to be a secure one, and the vulnerabilities of American democracy to informational attacks reflect problems in democratic institutions more than the unerring skill and craftiness of the attackers.

Addressing democracy's dilemma will involve figuring out new tradeoffs between openness and vulnerability, but working out such tradeoffs has always been a core problem of democracy. It also suggests that radical changes to make democracy fairer aren't simply more just in themselves, but are also justified by security concerns. That doesn't mean that such changes are easy to accomplish in a political system that seems almost purpose designed to stymie such reforms, but it does provide a broader rationale for them, as well as the beginnings of a systematic approach for thinking through what they involve.