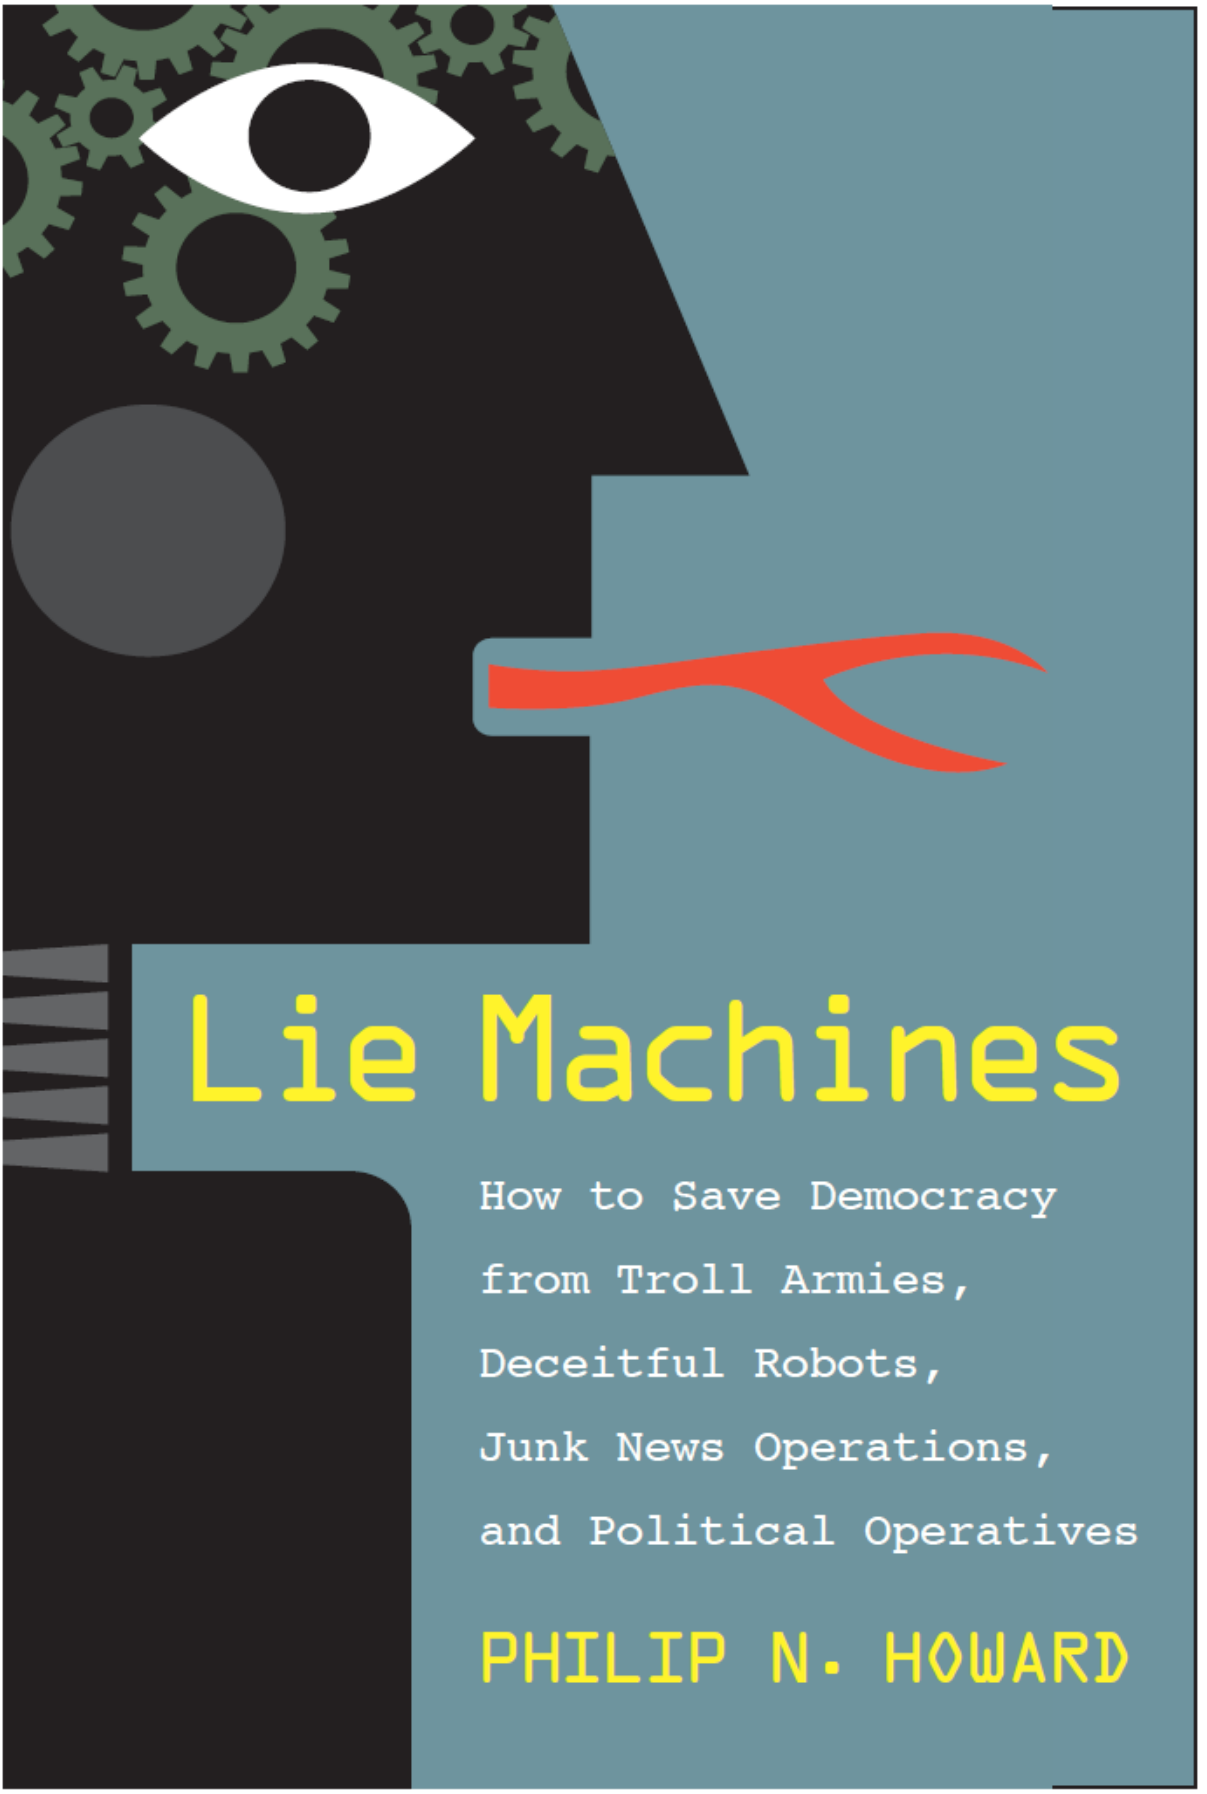

Lie Machines

How to Save Democracy from Troll Armies, Deceitful Robots, Junk News Operations, and Political Operatives

PHILIP N. HOWARD

1

The Science and Technology of Lie Machines

How is it possible to make convincing arguments, without evidence, about the causes of or solutions to important social problems? Lie machines do this work, and they seem to be more and more active in public life. These social and technical systems generate and spread political lies by subverting the people, organizations, and institutions in which we have most confidence. We tend to trust our instincts for evaluating the truth, and we tend to trust friends and family as sources of good information; as sources of news, many of us prefer social media to professional print media.[1]

Lie machines vary in size and scope. Some are large and permanent, others more liminal—temporary teams formed for a campaign or an election. But the inventory is growing, and in authoritarian regimes lie machines are often partnered with overt political surveillance initiatives. China employs two million people to write 448 million messages a year.[2] The previous Mexican president had seventy-five thousand highly automated accounts cheering for him and his policies.[3] Vietnam has trained

ten thousand students to deploy as cyber troops to combat online dissent.[4] The Thai government runs a special program rewarding children who report on any neighbors and family who protest using social media.[5] Huge bursts of automated activity occur whenever the presidents of the Philippines or Turkey score political victories.

Unfortunately, these massive mechanisms for social control don't only operate in authoritarian countries. In democracies, the surveillance is done by the social media firms themselves and not directly turned over to the government. However, the data gleaned by surveilling social media activity is still packaged and analyzed—it's just that the customers for this information are politicians, lobbyists, political consultants, and advertisers who are interested in what we're thinking about. Since 2016, every major democracy has suffered in some way: policy conversations go off the rails because of fake news; public understanding of critical issues is warped by well-advertised misinformation; surprising domestic political outcomes are shaped by hostile foreign actors.

Lie machines are large, complex mechanisms made up of people, organizations, and social media algorithms that generate theories to fit a few facts, while leaving you with a crazy conclusion easily undermined by accurate information. By manipulating data and algorithms in the service of a political agenda, the best lie machines generate false explanations that seem to fit the facts.

Lie machines have three main components, and if we examine the parts and understand how they work together, we can devise ways to take them apart or even prevent them from being built. The first component is the producer of lies that serve an ideology or the interests of political elites. These producers are often political candidates, campaign teams, political parties, and agencies of authoritarian governments, but

they can be marginal political actors who just want to disrupt public life. The second component is the distributor of lies: the algorithm that social media firms provide for spreading content. Social networks like Facebook, Instagram, Twitter, and WhatsApp—and the data they collect about our political preferences—are the distribution system for political lies. The third key component is the marketer of political lies, usually a consulting firm, lobbyist, or political hired gun who profit from selling, amplifying, and promoting disinformation. Big lie machines have production, distribution, and marketing systems that cross international borders, and in the chapters ahead we'll examine each component and evaluate how it works.

Large data sets have revolutionized politics by making it easier for politicians to understand what voters want. For decades politicians have used polling data and surveys to interpret what voters are thinking about. But traditional polling methods have always had limitations. And the polling data they gathered was either attitudinal (what people think) or aspirational (what people hope for). Now we have behavioral data about what people *actually do,* so lobbyists, campaign managers, and politicians can make much more powerful inferences about voters: Do you say you are pro-life on abortion issues but use your personal credit cards to buy contraceptives? What kinds of magazines do you subscribe to, and which online news sources do you spend the most time reading? How might your shopping choices reveal something about your thinking on environmental issues? All this data yields political information.

Lobbyists, political campaign managers, and politicians use social media to communicate directly with voters. Tools like Twitter and Facebook allow campaigners to communicate without worrying about how journalists and editors may change or interpret the campaign messages. And increasingly,

campaign managers employ automated tools to do this, with data generated by technology users—by us. Such automation allows a campaign to plan its communications far in advance, so that a greater volume of personalized, prefigured content can be sent to each voter even more quickly. Sometimes we are aware of the data trail we leave behind, but rarely can we see how it gets aggregated and merged with other data sets. And seldom do we see how it is made relational—put in context with data from our neighbors, friends, and families.

It can be hard to fully grasp where all that data comes from. Even though we are aware that our smartphones and computers keep track of our activities, we rarely think of the smart refrigerators, cars, lightbulbs, and other everyday objects that increasingly track our behaviors. And seldom do we see how this data is applied or passed to sophisticated artificial intelligence for analysis. Globally, each day, we generate five hundred million tweets, send 294 billion emails, put four petabytes of content on Facebook, send sixty-five billion messages on WhatsApp, and do five billion searches on Google. That's a lot of data about our thoughts, feelings, and activities. Yet there is even more data that connected devices around us collect—details that can also reveal much about our behavior. For example, every day each connected car generates four terabytes of data about how the engine is performing and where it is taking the driver.[6] Our cellphones and portable electronics generate similar data about performance, contents, and where we take them.

The Great Transformation

The internet has always been a powerful tool for political communication. But the tone and timbre of what was communicated, and for which purposes, has changed over time. In the

early days of the internet, when lobbyists wanted to push a piece of legislation or candidates ran for election, they could use the web to publish their platform, buy online banner ads, and trade messages with voters and supporters. This was still a one-to-many mode of communication, in that the political campaign would put content online and anyone could visit the webpage or consume the content if they were interested. More recently, social media platforms have introduced new mechanisms for political communication, including ways for users to involve their own networks of family and friends in clusters of political affiliation, content sharing, and data mining.

The people who follow your social media accounts are interested in what you say, and social media firms sell access to your entire network of family and friends. More importantly, social media platforms allow your followers to pass content to their own networks of family and friends. Campaign communications conducted over the internet can involve generic messages. But many messages that arrive on social media often look like personal communication because, by definition, we've included our friends and family in our social networks.

Political campaign managers have quickly made use of the networks of trust and reciprocity we reveal to social media firms. And many kinds of campaigners have taken advantage of network effects to spread appeals for support. During the 2008 US presidential election, Barack Obama's campaign organization used social media networks with great success. The campaign's technology strategy allowed supporters to network among themselves, offloading the organizational costs of campaigning onto energetic volunteers.[7] And during the Arab Spring protests of 2010–12, democracy advocates used social media networks to support one-to-one communication on a large scale to coordinate massive protests that brought down author-

itarian regimes in Tunisia, Egypt, Libya, and Yemen.[8] Regimes that didn't collapse were forced to make a range of political and financial concessions to appease popular demands. Having seen activists use social media as a mechanism for social mobilization, powerful lobbyists and authoritarian regimes began using the very same technologies for social control and to advance their causes.[9]

Microtargeting is the process of preparing and delivering customized messages to voters or consumers. Television, newspaper, and radio advertising can be targeted at broad categories of people. For example, men generally respond well to certain colors, keywords, pictures, and narratives, while women respond to others. Political campaign messages can be scripted to contain cultural cues that might attract or be familiar to, for example, urbanites, retirees, college students, or pregnant women. Campaign ads can be targeted using several demographic categories at once. As a concept, microtargeting involves preparing and delivering a message that is customized for particular individuals. With the right data, microtargeting can even design—in real time—ads customized to certain individuals. By assembling data from our personal credit card records, the national census, and our social media use, political campaigns can make some good guesses about our opinions on the issues of the day.

Some of the most effective lie machines, however, are built by political parties, authoritarian regimes, and radical social movements. Organized parties and movements attract supporters by using as much personal data on people as possible, whether targeting constituents, voters, members, consumers, or any other group of people. Sometimes this data is used for what I elsewhere term *political redlining*.[10] This is the process of deciding which people your campaign *doesn't* need to engage with. If, for example, people of color under a certain in-

come level rarely vote, or if they always vote for your opponent, it may not be worth spending time in their neighborhood trying to convince them to vote for you. If a city has consistently voted for one party for decades and there's no evidence that public sentiment is changing, why put campaign resources into advertising there? Of course, this is a rhetorical question, because it *is* healthy for civic life to have consistent, high-quality information and debate running in every neighborhood.

Bad Prospects

For the most part, when authoritarian governments, industry lobbyists, and shady politicians get hold of personal data, it is because they want to figure out how to build an effective campaign for winning support and legitimacy. Sometimes such data helps a politician or lobbyist better understand voters. When a crisis emerges or a major decision is needed, this data can allow political leaders to take the pulse of the community they serve. But usually personal data isn't used for substantive engagement; instead, it is applied for strategic reasons, to mobilize supporters and deepen existing commitments. Social media platforms are the delivery mechanism for the misinformation, for maneuvering or manipulating the individuals in the data set. In the chapters ahead, we'll see how many kinds of strategies, purposes, and outcomes there can be.

Most citizens don't know how much information about them is being used: it is bought and sold by campaigns, advertising firms, political consultants, and social media platforms. Sometimes political campaigns have clear sponsors and declare their interests, but at other times such motives are hidden. Sometimes the public can trace who is spending what resources on which ads during a campaign season, but at other

times the advertising is discrete and directed and not archived anywhere. Most democracies have independent elections regulators tasked with ensuring fair play, but not all such agencies are well funded. Sometimes campaigns deliberately spread misinformation, lies, and rumors about rivals and opponents that poison public life.

As far as I have observed over several years of investigation, political campaigns in the United States—and many other democracies—rarely give much thought to managing the data they collect. They don't have data management plans the way hospitals, government agencies, or even private firms do. Usually individual campaign managers and consultants maintain their own data sets, and this data travels with them from campaign to campaign. Political campaigns that are well resourced may commission surveys and political polls, and sometimes the polling firm will provide copies of the raw data, but usually the pollster writes up summary reports and interprets the trends for the campaign. Political campaigns that are very well resourced will hire consultants to build and maintain specialized political data sets that merge data from many sources: campaign contact records about voters, party membership information, public records, survey results, census data, social media data, and credit card data. Usually the consultants treat this information as their intellectual property and the political client doesn't have access to the raw data. Some consulting firms build trusted relationships with a political party so that strategic information that the party contributes stays among groups that share a political agenda.

Political campaigns have few systematic safeguards to protect data. As the 2016 Russian hack of the US Democratic Party, or the 2019 Russian hack of the UK Conservative Party illustrates, even the wealthiest political parties have trouble

keeping personal records secure. Some countries have privacy commissioners, but only since the Facebook–Cambridge Analytica scandal—in which millions of Facebook users' personal profiles were data mined and used for creating political ads—have countries started looking deeply into campaign practices around data. Most organizations and websites, such as social media platforms, credit card companies, and campaigns, have terms of service agreements. But usually these make a weak commitment to keeping data safe, saying that they will share data only with third parties they have evaluated for trustworthiness. In my experience, when you look at the list of third parties, you'll see quite a long list of subcontractors, individual consultants, and small ad agencies with their own inadequate data management practices.

If you want to explore what kind of political inferences can be made from the data you have been generating over the years, simply review your credit card records or get a copy of your credit history. First, know that this collection of data is only a fraction of what exists about you. It is sometimes called core data because along with demographics such as your age, gender, and race, it is a very basic set of facts about who you are, where you've been, and what you've been doing. It probably contains straightforward information about your address—probably all the addresses you've ever had—and that information is used for direct mailing campaigns. But core data usually also includes telephone number records so consultants can make robocalls or have third-party call centers contact you. Depending on the country you live in, there may be publicly available records of the political donations you have made or the political parties you belong to. When information about your political affiliations is married to your credit card records, a campaign professional has the data to make a whole lot

of inferences about how to organize the facts into a theory you are likely to believe.

Democracy Encoded

New technologies always inspire hot debate about the nature of politics and public life. When activists used Facebook to organize protests across North Africa, journalist Malcolm Gladwell argued that only face-to-face interaction could make a true social movement and a real revolution.[11] Technology writer Clay Shirky countered that it no longer made sense to distinguish street politics from internet politics.[12] Evgeny Morozov of the *New Republic* agreed that the internet was useful in politics, but mostly for dictators, internet activist Eli Pariser argued that too many of us were using technology to immerse ourselves in information bubbles that protect us from exposure to challenging new ideas.[13] I argued that everyone was focusing on the wrong internet—that we should be concerned about data from our own devices shaping politics, not content from our browsers.[14] This was the means by which political elites would "manage citizens" in the years ahead.[15]

Unsurprisingly, research has disproved some of these claims and confirmed others. Gladwell was wrong in that social media proved to be a sufficient cause of several contemporary political uprisings and revolutions.[16] Shirky was prescient in identifying the ways in which the internet supported new modes of organizing, modes that could catch dictators off guard.[17] Morozov was correct in suspecting that the bad guys would quickly learn to use new technologies in their organizational efforts and would work to catch democracy advocates and put them in jail, though his fatalism muddled his thinking. But as we'll see in the chapters ahead, it turns out that the

big, modern mechanisms for political manipulation depend on both our social media output and device-level data about our behavior: algorithms use such data to render the best models for doing targeted political advertising and organizing voter contact.

Despite all the debate about the impact of the internet on politics and public life over the past decade, no one thought that social media could be used to threaten so deeply the primary exercises of democracy: elections and referenda. Spreading vast quantities of political misinformation before voting day, in cleverly targeted ways, was not what social media was built for. Social media platforms were created from the ground up for advertising, but not to sell us political candidates. Yet as these platforms sought ways to monetize their assets—peoples' attention—they purposefully took advantage of our cognitive biases. More specifically, social media platforms served up our cognitive biases to political and ideological projects. At first, it was the toughest strongmen in Egypt, Turkey, and Pakistan who used social media as an instrument for social control. In democracies, far-right and far-left parties occasionally pushed the envelope with campaign tricks that either broke the law or violated the spirit of democratic discourse. What we didn't anticipate was the degree to which mainstream politicians in established democracies would use social media to manipulate their own publics.

In 2016, the campaign to have UK voters reject the European Union generated unusually bold lies about the costs and benefits of being an EU member. This campaign effectively used social media to poison debate and muddy issues. In the United States, pro-Trump bots, trolls, and ad buys greatly inflated a range of sensational, conspiracy, and extremist myths about Hillary Clinton. The Russian government and domestic

groups sharing political affinities sank significant resources into the distribution of polarizing content over Facebook and Twitter, expanding to include other social media platforms such as Instagram and YouTube in the years after the election of Donald Trump. Globally, many democratic elections since 2016 have suffered from some loss of credibility because of the operation of similar lie machines. In countries such as Brazil and the Philippines, populist leaders have embraced the same social media strategies used by authoritarian governments. Even in strong democracies like Canada and Germany, a few shady politicians get caught using these tools to suppress voter turnout, misinform voters, and undermine public confidence in electoral outcomes.

My research team at the University of Oxford has misinformation trends on social media in dozens of countries. During the presidential election of 2016 in the United States, there was a one-to-one ratio of junk news to professional news shared by voters over Twitter. In other words, for each link to a story produced by a professional news organization, there was another link to content that was extremist, sensationalist, or conspiratorial or to other forms of junk news. This is the highest level of junk news circulation in any of the countries we have studied.[18] Critically for politics in the United States, this misinformation was concentrated in the most competitive districts—where small shifts in popular opinion had big consequences.[19] Moreover, the audience for misinformation was primarily Republican voters who supported Trump.[20] Misinformation campaigns are often launched using highly automated accounts and fake users, and such accounts promoted significant amounts of content from Russian news sources, links to unverified material on WikiLeaks, and other forms of junk news. We have found that it is not just bot accounts that

spread junk news; at the right volume level, junk news can permeate deeply into networks of human users.[21]

Just a few months after the 2016 US elections, we demonstrated that disinformation about national security issues, including from Russian sources, was being targeted at US military personnel and veterans and their families.[22] During the president's State of the Union address, we learned that junk news is particularly appetizing for far-right white supremacists and President Trump's supporters (though not "small c" conservatives).[23] Some of this junk content originates in accounts managed by foreign governments.

Foreign intervention in politics, using social media, is ongoing, and it doesn't just involve the Russian government. As protests started rocking Hong Kong in 2019, Chinese government propagandists activated their social media networks to convince the world that the activists were violent radicals with no popular appeal. By 2020, seven countries were running misinformation campaigns targeting citizens in other countries: along with Russia and China, there were similar operations in India, Iran, Pakistan, Saudi Arabia, and Venezuela.[24] Whether built to target neighboring countries or domestic voters, these mechanisms do the same thing—they produce, disseminate, and market lies.

What Is a Lie Machine?

A *lie machine* is a system of people and technologies that distribute false messages in the service of a political agenda. The system involved can include many kinds of organizations, individuals, and relationships, from formal paid-employment relationships to those of volunteer associations and affinity groups who produce content and share it over networks of family and

friends. The people producing this have either a political or a financial incentive to do so—they work for political parties or they are political consultants and lobbyists for industry groups. But the most important people in these networks are citizens who make the mistake of passing a piece of misinformation from a junk news source, fake user, or political shill to broader networks of real citizens and voters. Some of the people involved in the mechanism are programmers who set up vast networks of highly automated accounts. These citizens have a disproportionate share of the public conversation because of the fake user accounts they control.

Social media platforms, search engines, and myriad devices provide the technology and infrastructure for delivering misinformation directly to citizens at key moments in a public conversation. Usually several technologies are used in a lie machine. For example, data from your credit card purchases may be used to create a digital political profile of you as a social media user. Then several of the platforms and websites you visit display political ads that are likely to trigger your interests and attention.

Whether through accidental sharing or purposeful distribution, an array of political actors and average citizens distribute and disseminate content by liking, retweeting, repurposing, and repackaging the content. Subtle alterations help the lies avoid spam filters, validity checks, and human editors and spread across social networks.

Unfortunately, there are many varieties of lies and what Caroline Jack calls "problematic information."[25] And the media's dependence on social media, analytics and metrics, sensationalism, novelty over newsworthiness, and clickbait gives the producers, distributors, and marketers of political lies the ability to project misinformation and disinformation to a wide audience.[26] *Misinformation* is contested information that re-

flects political disagreement and deviation from expert consensus, scientific knowledge, or lived experience. *Disinformation* is purposefully crafted and strategically placed information that deceives someone—tricks them—into believing a lie or taking action that serves someone else's political interests. *Junk news* is political news and information that is sensational, extremist, conspiratorial, severely biased, or commentary masked as news.

The messages, pictures, videos, and other content that lie machines distribute isn't simply false advertising. It can include very personal messages from platform users who are being paid by foreign governments to bait you to join an argument or spread a rumor. It can include highly customized posts about conspiracies. Paid advertising takes advantage of a platform's own means of generating a profit for itself, while direct messaging involves users who work for a political agenda.

The key distinction between misinformation and false advertising for commercial products is that the lies are in service of some political ideology. And generating social media content in the service of a party or an ideology increasingly takes a formal work arrangement, with teams under contract, storyboards and scripts, dedicated office space, and significant financial commitments. Sometimes individuals buy ads on their own or work as coordinated volunteers to drive a political figure off social media. But this work can also involve formal employment with agencies that coordinate large, agile teams that can be hired on short notice to support a political communication campaign.

Many kinds of political actors put technology in the service of power. In almost all countries, it is the far-right, ultraconservative groups and populist leaders that get caught using such tools early and aggressively. Only on rare occasions have we caught left-wing populists and middle-of-the-road politi-

cians putting these tools and techniques to work for a political agenda, and usually it's only after someone else has used a lie machine to attack them. Foreign governments meddle in the domestic affairs of neighboring countries, usually to sow local discord or target a politician whose ideas threaten the foreign government. When political parties do it, it's about projecting their values by stoking fear, uncertainty, and misunderstanding.

Lie machines are made up of parts—people and the technologies for disseminating the falsehoods they come up with. And when assembled properly, lie machines put information technology into service for a political ideology by generating computational propaganda. Deliberately using device networks, social media algorithms, and personal data to shape how we perceive the world is the work of regular advertisers but also of propagandists. Producing, distributing, and marketing propaganda is certainly an old craft, but these new, scaled-up mechanisms swiftly create and distribute content, with rapid testing and refinement, customized for individual consumption by that individual's personal and behavioral data.[27] This is not the broadcast era's propaganda of posters, government advertising, and leaflets dropped from airplanes. The modern lie machine is not a mass media instrument for broadcasting misleading information. It is a new, complex system of people and technologies with unique features.

The best lie machines involve components with three key functions: production, distribution, and marketing. The better these components work to promote a political lie, the greater the chances are of successfully manipulating public opinion on an issue. First, a well-running lie machine requires government agencies, political leaders, candidates for election, or political parties who produce misinformation in the service of their political agenda or big ideological project. Second, a lie machine needs a distribution system, which today takes the

form of the platforms provided by social media firms. The distribution system of bots, fake accounts, and easily exploited social media algorithms provides a technical infrastructure for packaging the lie and delivering it to your inbox. The third important component of a lie machine is the marketing work and this usually involves consultants and commercial agencies who polish and bundle the information into junk news stories. They analyze both big-picture market trends and data about you to get the marketing strategy right. Political consultants, lobbyists, and advertisers make big money by refining such marketing strategies. In sum, the key components are the people who produce political lies, the social media firms that distribute the lies, and the consulting firms that market the lies to you.

Simply defined, *computational propaganda* is the misleading news and information that is created in the service of political interests and algorithmically distributed across social media networks. It can involve networks of highly automated Twitter accounts or fake users on Facebook. Sometimes the people running these campaigns simply pay for advertising and exploit the services that social media companies offer to all kinds of advertisers. Often the political campaigns or foreign governments that generate computational propaganda take advantage of the design of platforms. So for dating apps like Tinder, the automated accounts or fake users will flirt and then talk politics. For YouTube, the goal is to repackage mainstream content or create dedicated channels to keep users in a narrow stream of content. Strategies for Reddit involve creating crazy new conspiracy theories and encouraging users to transport those theories to other platforms. Doing so usually means breaking terms of service agreements, violating community norms, and using platform *affordances* in a way the engineers didn't intend. Occasionally it means breaking elec-

tion law, privacy regulations, or consumer protection rules. But it happens anyway—sometimes because of inadequate enforcement of existing laws, regulations, and rules or inadequate public policy oversight. At times the firms themselves don't consistently enforce their own terms of service, because doing so would mean losing revenue or upsetting a group of users.

What Do Lie Machines Look Like?

Many people are questioning whether social media platforms are threatening democracy. Concentrated in just a few hands, these firms are aggregating large data sets about public and private life—including data on demographics and public attitudes and opinions. Making all this intelligence available for commercial use allows elite political actors to control the most valuable resource possible in a democracy: our attention. Information infrastructure is an invaluable asset to lobbyists seeking to pass legislation, foreign governments interested in controlling domestic conversations, and political campaign managers working to win an election.[28] While the internet has certainly opened new avenues for civic participation in political processes—inspiring hopes of a democratic reinvigoration—the parallel rise of big data analytics, black-box algorithms, and computational propaganda is raising significant concerns for policymakers worldwide. In many countries around the world, divisive social media campaigns have heightened ethnic tensions, revived nationalistic movements, intensified political conflict, and even resulted in political crises—while simultaneously weakening public trust in journalism, democratic institutions, and electoral outcomes.[29]

The first troll armies appeared in Russia in 2007, during Vladimir Putin's second term as president. I traveled there that summer to meet these early patriotic bloggers who went to

camps on the outskirts of Moscow to talk about the future of their country—and learn to code. After multiple successful campaigns of misinformation, the use of trolls spread as other governments, political parties, and lobbyists saw Russia's troll armies in action. By 2016, some twenty-five countries were spending upward of $2.5 billion a year on large cohorts of commentators, reviewers, harassers, and advocates who manipulated public life over social media. By 2020, more than seventy countries had organized social media misinformation teams.

These days, troll armies take many different organizational shapes. In the Philippines, President Rodrigo Duterte paid a small army of social media users not just to post upbeat messages about his candidacy for president but to threaten and harass his critics and opponents. In a country with one of the highest rates of extrajudicial killing and journalist murder rates in the world, a few threatening social media messages can be enough to drive many reasonable people out of public life. Given the great variety of lie machines in operation, what is the best way to frame the problem, explain it, and work out the best ways to break the lie machines down?

Lie Machines as Sociotechnical Systems

For many years, being a social scientist has meant purposefully gathering large amounts of information about people, their relationships, and their ideas. But the internet has also affected how we do such research, because it has become harder and harder to build meaningful explanations of social phenomena without making room for the affordances of information infrastructure. We used to think that people were the primary causes of social change. Many disciplines of social science inquiry investigate the role of media, technology, and information infrastructure and treat these things as tools

for individuals, networks, and communities. But increasingly we've had to admit that such a fundamental infrastructure as the internet both enables and limits our activities. Causal factors are often conjoined, and sensible explanations of modern politics need to give attention to the causal role of information technology.

A *sociotechnical system* is made up of people and their tools, and seeing both in interaction gives you a better analytical frame because you can observe how the material world provides both capacities for and constraints upon human action. We may all feel that we are agents, yet our ability to have an impact on the social and material world around us is largely determined by our access to information and our ability to communicate. Having a good education, consuming high-quality news, and using the latest information technology allows us to have an impact. These may give some of us the capacity to change the world, but for other people, not having access becomes a constraint.

Understanding the momentous changes and big challenges in how we do politics must mean understanding both people and technology. It is possible to be too technologically determinist, however. Arguing that the internet alone is responsible for social decay gets us nowhere. But it is also possible to be technologically underdetermined: we must also admit that the social media platforms we've built for ourselves have affordances that help us behave badly, circulate misinformation, cause other people harm, and degrade public life.

However, trying to understand computational propaganda only from a technical perspective—as hardware and algorithms—plays into the hands of those who use the technology to undermine democracy and the firms that profit when that happens.[30] The very act of describing something as

purely technical or in very mechanistic terms may make it seem unbiased, inevitable, or easily fixed by an engineer or a change in settings. This is clearly a dangerous position to take, and we must look to the emerging discipline of social data science to help us understand the complex sociotechnical issues at play and the influence of technology on politics. As part of this process, social data science researchers must maintain a critical stance toward the data they use and analyze to ensure that they are critiquing as they go about describing, predicting, or recommending changes in the way technology and politics evolve. If academic research on computational propaganda does not engage fully with the systems of power and knowledge that produce it—the human actors and motivations behind it—then the very possibility of improving the role of social media platforms in public life evaporates.[31] We can only hope to understand and respond appropriately to a problem like computational propaganda by undertaking computational research alongside qualitative investigation. This lets us see and understand the sociotechnical system.

Computational propaganda, then, is both a social and a technical phenomenon. It is produced by an assembly of social media platforms, people, algorithms, and big data tasked with the manipulation of public opinion.[32] *Computational* propaganda is of course a recent form of the propaganda that has existed in our political systems for millennia—communications that deliberately subvert symbols, appealing to our baser emotions and prejudices and bypassing rational thought to achieve the specific goals of its promoters—and is understood as propaganda that is created or disseminated by computational means. Automation, scalability, and anonymity are hallmarks of computational propaganda. The pernicious advantage of computational propaganda is that it enables the rapid

distribution of large amounts of content, sometimes personalized in order to fool users into thinking that messages originated in their extended network of family and friends. In this way, computational propaganda typically requires bots that automate content delivery, fake social media accounts that need little human curation, and junk news that is misinformation about politics and public life.[33]

Thus, the mechanism we need to understand—and rebuild—is one that involves both people and technology. People generate political misinformation; social media platforms such as Facebook, Instagram, and Twitter spread it around. Politicians generate negative messages; Twitter distributes the anger and outrage. Of course, journalists still generate high-quality journalism, and politicians do have positive, constructive ideas about public policy. But long-form, high-quality journalism and interesting new ideas about public policy aren't disseminated as widely as extremist, conspiratorial, and sensational junk news. If we analyze the entire mechanism, with its social and technical parts, can we redesign it to support democratic norms?

Taking the Lie Machines Apart

Public life is being torn apart. Lie machines sow distrust and infect political conversations with anger, moral outrage, and invective in ways that forestall consensus building. It is not simply that social media may have side effects, making us dependent on our screens for news and information, or that our mobile phones may be isolating us from our neighbors.[34] Troll armies, bot networks, and fake news operations are formal structures of misinformation, purposefully built. But to understand how they work, we must go behind the scenes of some of the most engrossing and appalling stories of modern

political intrigue and manipulation, explain the systems behind them, look ahead to how advanced technologies will be used, and give readers tools to break them down. And ultimately, analyzing these components isn't enough: we need to know how to fix the problem. So this book concludes by presenting ways that we can protect ourselves and help destroy the complex lie machines now threatening our democracies.

Most of the investigative journalism around this topic builds headlines around particular events involving trolls, bots, or fake news. But these components work very well together. Trolls launch negative campaigns, usually for radical movements or as paid advertising. They amplify their voice through automation and social media algorithms. Automated social media accounts have more impact when they link to fake news. Fake news has more sophistication and audience appeal when artificial intelligence is used to manipulate big data about public life. And very soon, the internet of things will generate the biggest data of all. The *internet of things* is made up of networks of manufactured goods with embedded power supplies, small sensors, and an address on the internet. Most of these networked devices are everyday items that are sending and receiving data about their conditions and our behavior. Such devices generate massive amounts of data, and if we are not careful, they could become part of truly enormous mechanisms for social control.

To understand properly how lie machines operate, we need lots of evidence. From interviews and fieldwork with political organizations and consultants to big data analysis of large amounts of content from Twitter and Facebook, we need to be able to explain how and why these assemblages get built if we are ever going to work out how to take them apart. So the analytical effort ahead involves evidence from different countries around the world, expert interviews, and exclusive, rare

data sets sourced from the social media companies themselves. I present original research on government agencies that are spending public money or retasking military units to keep certain hashtags trending, rile up particular groups of people, or defame rising political figures.

The recent slew of surprising political outcomes, strained civic dialogues, and polarizing debates seem—at face value—to be related to social media and the internet. In analyzing a political problem, our first instinct is often to try and identify the one person or factor to pin all the blame on. In the US context, one explanation is that the Russians are responsible for the current political climate because of their interference with how voters deliberated in 2016, 2018, and 2020.[35] A rival theory assigns blame to white supremacists and the far right for successfully using social media to radicalize mainstream conservatives.[36] But both explanations are incomplete because they miss the full perspective of the global political economy behind the lie machines that work on US voters.

How did we get here, and whom do we blame? Is Russia's Vladimir Putin behind it all? Is it our own fault for being addicted to Facebook and Twitter? Unfortunately, the answer is complex, involving both people and technologies in a sensible causal story about how big political lies get produced, distributed, and marketed to us.

The first components that we need to examine are the politicians and governments that generate big lies. As a starting point, we need to track the rise of organized teams of people dedicated to social media manipulation—groups of people often called *trolls.* The first troll armies, which appeared in Russia in 2007, are now organized by the Russian government's Internet Research Agency (IRA). I illustrate how they are used to actively spin government operations and take a close look

at the playbook used by social media trolls during a week in which Russia's ruling elites were accused of assassinating a democracy advocate.

The second core component to examine is the distribution method for political lies. We need to scrutinize the automation and algorithms that disseminate misleading political information. What are bots and who makes them? When do people believe a bot, and what impact do bots have on public opinion? I focus on the mechanics of automation on several platforms and on how the mechanism of automating political communication has become more and more sophisticated over the years.

Reset to lose a line

The final components of lie machines to understand are the consultants, marketers, and advertisers who put political lies into a distribution system—and profit by doing so. Many outrageous political stories, rumors, and accusations spread rapidly over social media, and there are businesses that profit by marketing, amplifying, and advertising political lies. In 2016, bots were successful in spreading a crazy story, often called #pizzagate, that supposedly linked Hillary Clinton with a pedophilia ring based out of a pizza parlor in Washington, DC. In 2020, it was automation on TikTok and Twitter that tried to convince local activists and the world at large to dismiss Hong Kong's democracy advocates as violent radicals. Every country now has similar kinds of politically potent lies—stories that remain believed long after they have been disproven. Who takes a potent piece of misinformation that serves the interests of political elites or some ideological agenda, does the market analysis, and unleashes a marketing campaign over social media? Who are the political operatives who buy and sell our data, make or break politicians, and distribute political lies over the internet?

After I've broken the machine down to these compo-

nents, the next task is to figure out how the parts fit together and affect public life. Understanding the impact and influence of a lie machine means tracing out the flow of information from lie producers through social media distribution systems and marketing plans by professional political consultants. We could call it a case study, archival research, process tracing, or a study of political economy, but it is simpler to say that we need to follow the money.

Lastly, if we can break a lie machine into component parts and then see how the lies and money flow, can we anticipate how such mechanisms might evolve in the years ahead? More importantly, can we catch and disassemble them permanently?

Unfortunately, bots are just early forms of automation that have filled our inboxes and social media feeds with junk. Political chatbots backed by sophisticated machine learning tools are now in development, and these automated accounts provide a much more humanlike, interactive experience. They are not simply scripted bots that can talk about politics. Are you sure we have evolved from primates? Does smoking cause all cancers or just some, and might this connection vary by gender and cigarette brand? There's a chatbot that will make you reconsider such things. Is climate change real? Should we inoculate our kids?

Chatbots have become the hot tool for industry lobbyists seeking to promote junk science. And the next step is to put *artificial intelligence* (AI) algorithms, which simulate learning, reasoning, and classification, behind these bots. Several militaries are looking at ways of using the latest nascent AI personalities and machine learning algorithms for political engagement. Some of the best stories are coming out of China, where closed media platforms allow for large-scale experimentation on public opinion. What can we say now about the political

biases in algorithms, and how will AI affect political participation in the democracies?

Most of us are used to experiencing the internet through the browser on our smartphone or computer. But as our cars, refrigerators, and thermostats are also increasingly connected to the internet, how will the data from devices with chips and sensors, wireless radios, and internet addresses be used? What can we learn about political opinion by playing with such big data? Moreover, how will others use this data to coerce us and manage public life? The most basic AI systems are already starting to use complex data to engage real people in political conversations. I highlight some of the amazing ways that this data is being used for meaningful political impact. When sophisticated AI gets to play with big sets of behavioral data from the network devices around us, someone somewhere will attempt to make political inferences from the analysis. Someone will always try to build better lie machines that offer ever more complete systems for understanding our behavior and pushing political opinion. I accept cynicism about the current impact of social media on democracy but not fatalism: civic action and policy leadership may well prevent the worst internet of things scenarios from becoming a reality.

There are multiple challenges before us if we want to live in functional democracies: A politician who doesn't like how a question is phrased dismisses the questioner as "alt-right" or "alt-left." A political leader who doesn't like how a news story is framed labels it "fake news." A political consultant who doesn't like the evidence comes up with "alternative facts." Growing numbers of citizens believe junk science about climate change and public health. Traditional pollsters can't call an election, and the surprising outcomes of elections seem to have their roots in manipulative leaders in other countries.

By closely examining lie machines, we can understand how to take them apart. I offer basic policy recommendations on how we can protect political speech while demolishing the mechanisms for producing, distributing, and marketing misinformation. I provide civic defense tips that should help us proactively protect ourselves in the years ahead. Yet the best way to solve collective problems is with collective action, so I also identify ways that our public agencies can protect us with policies that make it tough for these big lie machines to operate in our democracies. It is possible to block the production, dissemination, and marketing of big political lies, but we'll have to act together to do this effectively.