

A dark gray background featuring a complex network diagram. The diagram consists of numerous small, light gray circular nodes connected by thin, light gray lines. Several larger, medium-gray circular nodes are scattered throughout, acting as hubs or central points in the network. The overall effect is that of a dense, interconnected web of information or communication.

YOCHAI BENKLER,
ROBERT FARIS,
AND HAL ROBERTS

NETWORK PROPAGANDA

MANIPULATION,
DISINFORMATION, AND
RADICALIZATION
IN AMERICAN POLITICS

Network Propaganda

Network Propaganda

*Manipulation, Disinformation,
and Radicalization
in American Politics*



YOCHAI BENKLER
ROBERT FARIS
HAL ROBERTS

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Oxford University Press is a department of the University of Oxford. It furthers the University's objective of excellence in research, scholarship, and education by publishing worldwide. Oxford is a registered trademark of Oxford University Press in the UK and certain other countries.

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America.

© Yochai Benkler, Robert Faris, and Hal Roberts 2018

Some rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, for commercial purposes, without the prior permission in writing of Oxford University Press, or as expressly permitted by law, by licence or under terms agreed with the appropriate reprographics rights organization.



This is an open access publication, available online and distributed under the terms of a Creative Commons Attribution – Non Commercial – No Derivatives 4.0 International licence (CC BY-NC-ND 4.0), a copy of which is available at <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Enquiries concerning reproduction outside the scope of this licence should be sent to the Rights Department, Oxford University Press, at the address above.

Library of Congress Cataloging-in-Publication Data

Names: Benkler, Yochai, author. | Faris, Robert, author. | Roberts, Hal (Harold) author.
Title: Network propaganda : manipulation, disinformation, and radicalization in American politics / Yochai Benkler, Robert Faris, Hal Roberts.
Description: New York, NY : Oxford University Press, 2018. | Includes bibliographical references and index.
Identifiers: LCCN 2018020121 | ISBN 9780190923624 ((hardback) : alk. paper) | ISBN 9780190923631 ((pbk.) : alk. paper)
Subjects: LCSH: Presidents—United States—Election—2016. | Communication in politics—United States. | Political campaigns—United States. | Mass media—Political aspects—United States. | Social media—Political aspects—United States. | Internet in political campaigns—United States. | Disinformation—United States—History—21st century. | Radicalism—United States. | Political culture—United States. | United States—Politics and government—2009–2017. | United States—Politics and government—2017–
Classification: LCC JK526 2016 .B46 2018 | DDC 324.973/0932—dc23
LC record available at <https://lccn.loc.gov/2018020121>

1 3 5 7 9 8 6 4 2

Paperback printed by Sheridan Books, Inc., United States of America
Hardback printed by Bridgeport National Bindery, Inc., United States of America

Note to Readers

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is based upon sources believed to be accurate and reliable and is intended to be current as of the time it was written. It is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional person should be sought. Also, to confirm that the information has not been affected or changed by recent developments, traditional legal research techniques should be used, including checking primary sources where appropriate.

(Based on the Declaration of Principles jointly adopted by a Committee of the American Bar Association and a Committee of Publishers and Associations.)

**You may order this or any other Oxford University Press publication
by visiting the Oxford University Press website at www.oup.com.**

Contents

<i>Acknowledgments</i>	vii
------------------------	-----

PART ONE: Mapping Disorder

1. Epistemic Crisis	3
2. The Architecture of Our Discontent	45
3. The Propaganda Feedback Loop	75

PART TWO: Dynamics of Network Propaganda

4. Immigration and Islamophobia: Breitbart and the Trump Party	105
5. The Fox Diet	145
6. Mainstream Media Failure Modes and Self-Healing in a Propaganda-Rich Environment	189

PART THREE: The Usual Suspects

7. The Propaganda Pipeline: Hacking the Core from the Periphery	225
8. Are the Russians Coming?	235
9. Mammon's Algorithm: Marketing, Manipulation, and Clickbait on Facebook	269

PART FOUR: *Can Democracy Survive the Internet?*

10. Polarization in American Politics	295
11. The Origins of Asymmetry	311
12. Can the Internet Survive Democracy?	341
13. What Can Men Do Against Such Reckless Hate?	351
14. Conclusion	381
<i>Notes</i>	389
<i>Index</i>	443

Acknowledgments

THIS BOOK REPRESENTS an account and update from what has been a long intellectual journey for each of us. For well over a decade, we have collected and used data to study, understand, and describe the impact of newly emerging digital communication on society, politics, and democracy. We have not undertaken this journey alone, and this book has benefited from the input and support of countless people along the way.

We would first like to thank Ethan Zuckerman for his decade-long partnership with us to support both the intellectual work of this book and the development of the Media Cloud platform that enabled the core analysis in this book. A decade ago, we began developing the technical infrastructure for the data analysis platform which would eventually take the name Media Cloud. Spurred on by debates within the Berkman Klein Center, where Zuckerman, now Director of the Center for Civic Media at MIT, co-founder and co-Principal Investigator of the Media Cloud project, was a fellow and senior researcher, and across the broader academic community, we sought to develop better tools to empirically study the structure and function of digital media. At that time, the open web was the core of digital communication. Much of our attention was directed at studying the impact of blogs on public discourse while Facebook, Twitter, and YouTube were relatively recent additions to the digital landscape. While we debated whether blogging would democratize media production and strengthen, we set about the many years' process of building the Media Cloud platform to collect, parse, and analyze digital media.

This book is the result of many months of effort and has only come to be as a result of the generous input of our friends and colleagues from the Berkman Klein Center and beyond. We are especially indebted to our colleagues Nikki Bourassa, Bruce Etling, and Justin Clark, who have made important substantive contributions to this study, supported the overall

research enterprise, conducted analysis, gathered data, provided input and feedback on this book, and shaped our understanding of the issues. Kira Tebbe provided crucial assistance in the final editing and production of the book. Rebekah Heacock Jones helped get this research off the ground with research into political discourse on Twitter. Daniel Dennis Jones worked tirelessly in the production and publication of this work. We benefited from the insights and efforts of Zach Wehrwein and Devin Gaffney, who helps us to track and understand the propagation of frames and narratives from Reddit. Brendan Roach and Michael Jasper provided invaluable research assistance. Jonas Kaiser and Paola Villarreal expanded our thinking around methods and interpretation. Alicia Solow-Niederman worked tirelessly to debug early versions of the analytical methods that were used in this book. Urs Gasser and Jonathan Zittrain have extended valuable support that has enabled us to maintain this research for many years. John Palfrey and Colin Maclay provided critical institutional support for Media Cloud in its early stages.

We are grateful to our friends and collaborators at the Center for Civic Media at the MIT Media Lab who have worked with us in the development of the Media Cloud platform and contributed to the applied research it has supported. This work has built upon and fostered an unusually close and productive collaboration between our two academic centers. Rahul Bhargava, Linas Valiukas, and Cindy Bishop have helped to extend and translate the ideas and concepts of a large-scale media analysis platform into the current functionality of the Media Cloud platform upon which this research relied. Fernando Bermejo has been a valuable supporter and contributor to our collective work in this field. Natalie Gyenes and Anushka Shah provided research insights and Media Cloud expertise.

This research has also benefited from contributions of many outside the Berkman Klein community. John Kelly and Vlad Barash provided important insights into the role of social media in the election, leading us to new hypotheses and ideas that shaped the book's development. Matt Higgins helped lay a firm foundation of thought and hypotheses upon which this work was completed. Philipp Nowak provided valuable early research assistance. Participants of Data & Society's Propaganda & Media Manipulation Workshop in May 2017 provided valuable feedback and critical cross-examination that helped steer our earlier work toward this final version. We are also indebted to our editors at the Oxford University Press, Alex Flach and Emma Taylor, without whose initiative and support we would not have translated and extended our research into this book.

The research on the post-election period and additional research necessary to understand the institutional foundations we describe here as well as the production of the book benefited from the support of the Ford Foundation and the Ethics and Governance of Artificial Intelligence Fund. The original study of the election period, upon which significant portions of this book are based, was funded by the Open Society Foundations U.S. Programs. This work would not have been feasible without the investments made by a set of funders who have funded Media Cloud development over the years, including the Bill and Melinda Gates Foundation, the Robert Wood Johnson Foundation, the Ford Foundation, the John D. and Catherine T. MacArthur Foundation, and the Open Society Foundations.

We are also grateful to those who have invested time and resources into developing open tools and data for public interest research. We drew upon TV Archive data collected and made available by the Internet Archive, and used the tools and interfaces developed and made publicly available by the GDELT Project. Gephi, open source network analysis and visualization software, served as the engine for our network analysis.

PART ONE

Mapping Disorder

Epistemic Crisis

As a result of psychological research, coupled with the modern means of communication, the practice of democracy has turned a corner. A revolution is taking place, infinitely more significant than any shifting of economic power

WALTER LIPPMANN, *Public Opinion*, 1922

ON SUNDAY, DECEMBER 4, 2016, a young man carrying an assault rifle walked into Comet Pizza in Northwest Washington, D.C., to investigate reports that Hillary Clinton and her campaign chief were running a pedophilia ring from the basement of the pizza parlor.¹ A week later, a YouGov poll found that, however whacky the story, the young man was not alone in believing it; nearly half of Trump voters polled “gave some credence” to the rumors.²

Two weeks earlier, BuzzFeed’s Craig Silverman had published an article that launched the term “fake news.”³ Silverman’s article examined engagements with news stories on Facebook through shares, reactions, and comments and argued that the best-performing stories produced by political clickbait sites masquerading as actual news sites, often located offshore, generated more Facebook engagements than the top stories of legitimate news sites. On January 6, 2017, the Office of the Director of National Intelligence released a report that blamed Russia of running a disinformation campaign aimed to influence the U.S. election with the aim of helping Donald Trump get elected.⁴

The steady flow of stories reinforced a perception that the 2016 election had involved an unusual degree of misleading information flowing in the American media ecosystem. From claims during the primary that Jeb Bush had “close Nazi ties,”⁵ through claims during the general election that Hillary Clinton’s campaign was 20 percent funded by the Saudi royal family,⁶ the campaign was littered with misleading stories, often from sources that masked their identity or affiliation. Moreover, just as with the alleged pedophilia case,

Network Propaganda. Yochai Benkler, Robert Faris, and Hal Roberts.

© Yochai Benkler, Robert Faris, and Hal Roberts 2018. Published 2018 by Oxford University Press.

many of the stories seemed designed to elicit fear and disgust, as the titles of Breitbart's most widely shared stories on immigration exhibit: "Six Diseases Return to US as Migration Advocates Celebrate World Refugee Day," and "More than 347,000 Convicted Criminal Immigrants At Large in U.S."⁷

The 2016 U.S. presidential election won by Donald Trump followed closely on the heels of the equally shocking success of the Leave campaign in Britain's vote to exit the European Union. Both seemed to mark an epistemic crisis in contemporary democratic societies. As 2016 was drawing to a close, many in the United States and the European Union saw these events as signals that democracy itself was in crisis, buckling under the pressure of technological processes that had overwhelmed our collective capacity to tell truth from falsehood and reason from its absence. Brexit and the rise of far-right parties in countries such as France, Hungary, Austria, and even Sweden signaled a deep crisis in the pluralist, cosmopolitan, democratic project that was at the heart of the project of Europe. The victory of Donald Trump marked a triumph of a radical populist right-wing politics that had long simmered on the margins of the American right and the Republican Party: from the segregationist third-party candidacy of George Wallace in 1968, through Pat Buchanan's primary runs in 1992 and 1996, to the rise of the Tea Party after 2008. These remarkable political victories for what were once marginal ideologies appeared at the same time that democracies around the world, from the Philippines, through India, to Turkey saw shifts from more liberal democratic forms to a new model of illiberal, and in some cases authoritarian, majoritarianism.

Something fundamental was happening to threaten democracy, and our collective eye fell on the novel and rapidly changing—technology. Technological processes beyond the control of any person or country—the convergence of social media, algorithmic news curation, bots, artificial intelligence, and big data analysis—were creating echo chambers that reinforced our biases, were removing indicia of trustworthiness, and were generally overwhelming our capacity to make sense of the world, and with it our capacity to govern ourselves as reasonable democracies.

The first year of the Trump presidency brought no relief. The president himself adopted the term "fake news" to describe all news that was critical or embarrassing. By the end of his first year in office, the president was handing out "Fake News Awards" to his critics, and four in ten Republicans responded that they "considered accurate news stories that cast a politician or political group in a negative light to be fake news."⁸ While trust in news media declined in a broad range of countries, the patterns of trust and mistrust

differed widely across different countries. Together with Hungary and Israel, two other democracies with powerful right-wing parties, the United States was an outlier: distrust was high on average but markedly higher for one party affiliation.⁹

Echo chambers ringing with false news make democracies ungovernable. We can imagine a pluralist democracy in which populations contested elections and won or lost based on their votes, without ever sharing a viewpoint on what is going on in the world. Partisan press hurling accusations at the other party was, after all, the norm in nineteenth-century America. One party might believe that we are under attack from zombies and vote to counter this existential menace, while another party might believe that we are threatened by a long-term decline in productivity growth and vote to focus on that problem. Whoever won would design policies to counter what they saw as the major policy question of our times. The role of pluralist democracy would be to govern the rules of orderly transition from the zombie slayers to the productivity wonks and back with the ebb and flow of electoral success.

In practice, given the intensity of the “zombie-threat” party’s sense of impending doom, such a pluralist democracy would be deeply unstable. Some shared means of defining what facts or beliefs are off the wall and what are plausibly open to reasoned debate is necessary to maintain a democracy. The twentieth century in particular saw the development of a range of institutional and cultural practices designed to create a shared order out of the increasingly complex and interconnected world in which citizens were forced to address a world beyond their local communities, values, and beliefs. The medical profession, for example, rapidly and fundamentally transformed itself after the discovery of germ theory in 1876. Between 1900 and 1910, the American Medical Association grew from 8,400 to 70,000 members. This growth represented the transition from dubiously effective local medical practices to a nationally organized profession acting as an institutional gatekeeper for scientifically-based practices. The same pattern can be found in the establishment of other truth-seeking professions during the early twentieth century, including education, law, and academia. All of these professions organized themselves into their modern national, institutional forms in roughly the first 20 years of the twentieth century. Those years saw the emergence of, among others, the American Bar Association; the National Education Association; and the American Historical, American Economic, American Statistical, and American Political Science Associations.¹⁰

During this same critical period, journalism experienced its own transformation, into an institutionalized profession that adopted practices we

would recognize as modern objective journalism. By 1912 Columbia University's journalism school had been founded, which helped to institutionalize through professional training a set of practices that had developed in the late nineteenth and early twentieth centuries and that we now associate with objective journalism—detachment, nonpartisanship, the inverted pyramid writing style, facticity, and balance. Before this development, none of these attributes were broadly present in journalism.¹¹ These shifts in the professions in general, and in journalism in particular, were in turn part of the broad shift associated with modernism, employing rational planning, expertise, and objective evidence in both private sector management and public administration.

Since the end of World War II this trend toward institutionalized professions for truth seeking has accelerated. Government statistics agencies; science and academic investigations; law and the legal profession; and journalism developed increasingly rationalized and formalized solutions to the problem of how societies made up of diverse populations with diverse and conflicting political views can nonetheless form a shared sense of what is going on in the world.¹² As the quip usually attributed to Daniel Patrick Moynihan put it, "Everyone is entitled to his own opinion, but not his own facts." Politics was always centrally about identity and belonging and meaning, but in the decades following World War II, democracy operated within constraints with regard to a shared set of institutional statements about reality. Zombie invasions were out.

Zombie invasions are definitely back in. The year following the 2016 U.S. presidential election saw publication of reports¹³ and academic papers¹⁴ seeking to categorize the confusion, defining misinformation (the unintentional spread of false beliefs) and disinformation and propaganda (the intentional manipulation of beliefs), identifying their sources,¹⁵ and studying the dynamics by which they spread.¹⁶ This flurry of work exhibited a broad sense that as a public we have lost our capacity to agree on shared modes of validation as to what is going on and what is just plain whacky. The perceived threats to our very capacity to tell truth from convenient political fiction, if true, strike at the very foundations of democratic society. But it is important to recognize that for all the anxiety, not to say panic, about disinformation through social media, we do not yet have anything approaching a scientific consensus on what exactly happened, who were the primary sources of disinformation, what were its primary vectors or channels, and how it affected the outcome of the election. In this book we try to advance that diagnosis

by applying a wide range of tools to very large data sets and reviewing the literature that developed over the first year after the election.

The critical thing to understand as you read this book is that the epochal change reflected by the 2016 election and the first year of the Trump presidency was not that Republicans beat Democrats despite having a demonstrably less qualified candidate. The critical change was that in 2016 the party of Ronald Reagan and the two presidents Bush was defeated by the party of Donald Trump, Breitbart, and billionaire Robert Mercer. As our data show, in 2017 Fox News joined the victors in launching sustained attacks on core pillars of the Party of Reagan—free trade and a relatively open immigration policy, and, most directly, the national security establishment and law enforcement when these threatened President Trump himself. Our work helps to explain how a media ecosystem that initially helped the GOP gain and retain power ultimately spun out of control. From the nomination of Roy Moore as Republican candidate for the Alabama special Senate election over the objections of Republican Party leadership to Republican congressman Francis Rooney’s call to “purge the FBI,” and from the retirement of Paul Ryan from his position as Speaker of the House to evangelical leader Franklin Graham’s shrug at Donald Trump’s marital infidelities, a range of apparently incongruous political stories can be understood as elements of this basic conflict between Trumpism and Reaganism over control of the Republican Party. In the 2016 election, once the Trump Party took over the Republican Party, many Republicans chose to support the party that had long anchored their political identity, even if they did not love the candidate at the top of the ticket. Indeed, it is likely that the vehemence of the attacks on Hillary Clinton that we document in Chapters 3, 4, 6, and 7 were intended precisely to reduce that dissonance, and to make that bitter medicine go down more easily. Our observations, and the propaganda feedback loop we identify in Chapters 2 and 3 help explain both how such a radicalization could have succeeded within the Republican Party, and how that transformation could achieve an electoral victory in a two-party system that should, according to the standard median voter models favored in political science, have led the party rebels to electoral defeat and swept them into the dustbin of history. We leave until Part Three our historical explanation for how and when that propaganda feedback loop established itself in the right wing of American politics.

The bulk of this book comprises detailed analyses of large data sets, case studies of the emergence of broad frames and particular narratives, and synthesis with the work of others who have tried to make sense of what

happened at both abstract and concrete levels. Our goal is to understand which actors were responsible for this transformation of the American public sphere, and how this new public sphere operated through those actors so as to make it so vulnerable to disinformation, propaganda, and just sheer bullshit. Our heavy focus on data is complemented by an effort to make sense of what we see today in historical context, both political and cultural.

We take a political economy view of technology, suggesting that the fundamental mistake of “the internet polarizes” narrative is that it adopts too naïve a view of how technology works and understates the degree to which institutions, culture, and politics shape technological adoption and diffusion patterns. These, we think, were the prime movers of the architecture of American political media, and it is this finding that makes this book, for all its detailed focus on American politics and media, a useful guide for other countries as well. We argue that it would simply be a mistake for countries such as, say, Germany, to look at elections in the United States or the United Kingdom, see the concerns over online information pollution or propaganda, and conclude that the technology, which they too use, is the source of disruption. Different political systems, coming from different historical trajectories and institutional traditions, will likely exhibit different effects of the same basic technological affordances. So it was with mass circulation presses, movies, radio, and television, and so it is with the internet and social media. Each country’s institutions, media ecosystems, and political culture will interact to influence the relative significance of the internet’s democratizing affordances relative to its authoritarian and nihilistic affordances. What our analysis of the American system offers others is a method, an approach to observing empirically what in fact is happening in a country’s political media ecosystem, and a framework for understanding why the particular new technological affordances may develop differently in one country than another.

Dramatis Personae

Media and academic discussions of the post-truth moment have identified a set of actors and technological drivers as the prime suspects in causing the present state of information disorder, such as fake news purveyors, Russians, and so forth. These discussions have also employed a broad range of definitions of the problem. Before turning to our analysis, we offer, first, the list of actors who have been described as potentially responsible for disrupting American political communications, and second, precise definitions of the terms we will

use in describing the sources and forms of misperceptions that spread through the American media ecosystem.

“Fake News” Entrepreneurs/Political Clickbait Fabricators.—Before Donald Trump appropriated the term, the “fake news” phrase took off in the wake of Craig Silverman’s reporting on BuzzFeed about the success of fake election news stories.¹⁷ This reporting built on Silverman’s earlier story describing over 100 pro-Trump websites run from a single town in the former Yugoslav Republic of Macedonia. The Macedonian teenagers responsible had little interest in American politics but had found that by imitating actual news sites, and pushing outlandish hyperpartisan stories, they got lots of Facebook engagements from Trump supporters, which translated into very real advertising dollars.¹⁸ For a while, these websites received a lot of media attention.¹⁹ Their operators had figured out how to leverage a core affordance of Facebook—its ability not only to connect publishers with audiences, but also to generate revenues and distribute them to publishers able to elicit “engagements” on the platform. The social media entrepreneurs who created these sites were the perfect target of anxiety for traditional media: they diverted attention and advertising dollars from “legitimate” media, they manipulated Facebook’s algorithm, they were mostly foreign in these stories, and they were purely in it for the money. Here, we call them “clickbait fabricators,” and primarily address their role in Chapter 9. By “clickbait” we mean media items designed to trigger an affective response from a user that leads them to click on the item—be it an image, a video, or a headline—because the click itself generates revenue for the clickbait purveyor. While this can easily apply to many news headlines and much of online advertising, “clickbait fabricators” are individuals or firms whose product is in effect purely the clickbait item, rather than any meaningful underlying news or product. We use the “fake news” moniker to introduce them here because it was used early on to identify this particular threat of pollution from political clickbait fabricators. Elsewhere, we avoid the term itself because it is too vague as a category of analysis and its meaning quickly eroded soon after it was first introduced.

Russian Hackers, Bots, and Sockpuppets—Claims of Russian intervention in the U.S. election surfaced immediately after the hacking of the Democratic National Committee (DNC) email server, in June 2016.²⁰ By the end of the year, it had become an official assessment of the U.S. intelligence community.²¹ Over the course of 2017 and 2018 this set of concerns has been the most politically important, not least because of the criminal investigation

into alleged connections between the Trump campaign and Russia. Reports and documentation released by congressional committees shone particular attention on Russian propaganda use of Facebook advertising. Facebook itself, and later Twitter, issued reports confirming that they had identified instances of Russian interference. A range of independent academic and nonprofit reports confirmed the effort. The types of interventions described included the email hacks themselves—primarily the DNC and John Podesta emails—which provided grist for the partisan mill in the months before the election; and the use on Facebook and Twitter of automated accounts (“bots”), and “fake” accounts masquerading as something other than Russian agents (“sockpuppets”), which incited people on both the right and the left to protest, and pushed and gave particular prominence to anti-Clinton and pro-Trump messages. We dedicate Chapter 8 to assessing the Russian threat in detail.

The Facebook News Feed algorithm and Online Echo Chambers—A third major suspect was centered on the Facebook News Feed algorithm, although it extended to other social media and the internet more generally as well. To some extent, this was simply a reprise of the nearly 20-year-old concern that personalization of news, “the Daily Me,” would drive us into “echo chambers” or “filter bubbles.” To some extent it reflected a wave of newer literature concerned in general with algorithmic governance, or the replacement of human, legible, and accountable judgments with “black box” algorithms.²² In particular it reflected the application of this literature to politics in the wake of a series of experiments published by Facebook research teams on the News Feed algorithm’s ability to affect attitudes and bring out the vote.²³ It was this algorithm that rewarded the clickbait sites circulating the hyperpartisan bullshit. It was this algorithm that reinforced patterns of sharing in tightly clustered communities that supported the relative insularity of user communities. As a result, many of the most visible reform efforts in 2017 and 2018 were focused on revisions of the Facebook News Feed algorithm to constrain the dissemination of political clickbait and Russian propaganda. As with the case of the Russians, concern over the Facebook News Feed algorithm in particular, and over algorithmic shaping of reading and viewing habits in general, is legitimate and serious. In our observations, Facebook appears to be a more polluted information environment than Twitter or the open web. In Chapters 2, 3, and 9, we show that sites that are particularly prominent on Facebook but not on Twitter or the open web tend to be more prone to false content and hyperpartisan bullshit, on both sides of the political divide,

although there is more than enough pollution on these other media as well. But, we will explain why manipulations of Facebook's platform, like Russian intervention, were nonetheless not the primary driver of disinformation and confusion.

Fake news entrepreneurs, Russians, the Facebook algorithm, and online echo chambers provide normatively unproblematic, nonpartisan explanations to the current epistemic crisis. For all of these actors, the strong emphasis on technology suggests a novel challenge that our normal systems do not know how to handle but that can be addressed in a nonpartisan manner. Moreover, focusing on "fake news" from foreign sources and on Russian efforts to intervene places the blame onto foreigners with no legitimate stake in our democracy. Both liberal political theory and professional journalism consistently seek neutral justifications for democratic institutions, so visibly nonpartisan explanations such as these have enormous attraction. The rest of the actors, described below, lack this nonpartisan characteristic.

Cambridge Analytica—Another commonly blamed actor is the Trump campaign's use of Cambridge Analytica to manipulate behavior using artificial intelligence (AI)-driven social media advertising. The extent to which Cambridge Analytica, a U.K.-based data analytics political consultancy that had used tens of millions of Facebook profiles to develop techniques for manipulating voters, in fact used psychographic data and manipulated targets is debatable. What is clear is that the social media companies, Facebook in particular, helped the Trump campaign, as they would any paying customer, to use their deep data and behavioral insights to target advertising.²⁴ It is less clear, however, that there is anything wrong, from the perspective of American norms of electoral politics, with this campaign usage of cutting-edge, data-driven behavioral marketing. In 2012, when the Obama campaign used then-state-of-the-art data-driven targeting, post-campaign analyses feted the campaign geeks.²⁵ If there is a problem here, it is part of a much broader and deeper critique of behavioral marketing more generally, and how it undermines consumer and citizen sovereignty. We outline some of the events and the broader concerns in Chapter 9, explain why the threat is likely more remote than news coverage of Cambridge Analytica implied, and suggest how some of the proposed solutions may, or may not, help with this long-term threat in Chapter 13.

White Supremacist and Alt-Right Trolls—One of the most troubling aspects of the 2016 election and the politics of 2017 was the rise of white supremacists

in American politics. As Alice Marwick and Rebecca Lewis carefully documented, white supremacists, neo-Nazis, and other long-standing denizens of the American far-right found fellow travelers in young, net-native subcultures on Reddit, 4chan, and 8chan, graduates of the Gamergate controversy, and other online trolls, to undertake a meme war.²⁶ The core argument is that these decentralized, politically mobilized, and meme-savvy activists deployed a set of disinformation memes and framings that altered the election. Serious anthropological and computational work, in addition to the work of Marwick and Lewis, supports the argument that these meme campaigns had significant impact on the campaign.²⁷ Our own work detailed in the following chapters, however, aligns with that of researchers, including Whitney Phillips, Jessica Breyer, and Gabriella Coleman,²⁸ who were more skeptical of the central role assigned to “alt-right” online activists by some. In Chapter 4 we document how isolated the white supremacist sites were from the overall Islamophobic framing of immigration that typified right-wing media. In Chapter 7 we document how these activists intersected with Russian propagandists to propel stories up the propaganda pipeline, but also suggest that these events were, in the scheme of things, of secondary importance.

The impact of the white supremacists matters a great deal, because fear over their impact has created nettlesome problems for Americans concerned with democracy and the First Amendment; and for Europeans concerned with far-right propaganda on one hand, and the fear of American companies imposing their speech standards on Europeans on the other hand. Far-right activist meme wars undoubtedly represent core political speech, by a politically mobilized minority. It is hard to think of a clearer case for First Amendment protection. But many of the techniques involved in these campaigns involve releasing embarrassing documents, hateful drowning-out of opponents, and other substantial personal offenses. The substantive abhorrence of explicitly racist and misogynistic views and the genuine concern with the effects of the intimidation and silencing campaigns have increased calls for online censorship by privately owned platforms. The most visible results of these calls were the decisions by GoDaddy, Google, and Cloudflare to deny services to the Daily Stormer, a neo-Nazi site, in the wake of the white supremacist demonstrations in Charlottesville, Virginia, in the middle of 2017. In Europe explicitly Nazi content is an easier constitutional case, but questions of what counts as illegal and worthy of removal will remain central. A German law called the ‘NetzDG’ law, effectively enforced since January 2018, became the most aggressive effort by a liberal democracy to require online platforms to

police their systems. Aimed at hate speech in particular, the law imposed very large fines on major online platforms if they failed to remove speech that violates a broad set of German criminal prohibitions, some of which applied to much broader and vaguer categories than obvious hate speech. We offer a more detailed description of this law and its limitations in Chapter 13. That law will undoubtedly inform other countries in Europe and elsewhere as they decide to create their own versions of laws that push private platforms to impose what some would call “editorial control” and others “private censorship.” Our data support the more reticent approach, based on the scarcity of evidence of transformational impact of these extremists on the U.S. media ecosystem. Throughout our case studies we observe instances of alt-right memes trickling through the media ecosystem, but to do so they rely overwhelmingly on transmission by the more prominent nodes in the right-wing media network. These major right-wing outlets, in turn, are adept at producing their own conspiracy theories and defamation campaigns, and do not depend on decentralized networks of Redditors to write their materials. Given the secondary and dependent role that these sites have on the shape of the American media ecosystem, the gains from silencing the more insulated far-right forums may be less significant than would justify expansion of the powers of private censorship by already powerful online platforms in relatively concentrated markets.

Right-Wing Media Ecosystem—Our own contribution to debates about the 2016 election was to shine a light on the right-wing media ecosystem itself as the primary culprit in sowing confusion and distrust in the broader American media ecosystem. In the first two parts of this book we continue that work by documenting how the right-wing media ecosystem differs categorically from the rest of the media environment and how much more susceptible it has been to disinformation, lies, and half-truths. In short, we find that the influence in the right-wing media ecosystem, whether judged by hyperlinks, Twitter sharing, or Facebook sharing, is both highly skewed to the far right and highly insulated from other segments of the network, from center-right (which is nearly nonexistent) through the far left. We did not come to this work looking for a partisan-skewed explanation. As we began to analyze the millions of online stories, tweets, and Facebook sharing data points, the pattern that emerged was clear. Our own earlier work, which analyzed specific campaigns around intellectual property law and found that right and left online media collaborated, made us skeptical of our initial observations, but these proved highly resilient to a wide range of specifications and robustness

checks. Something very different was happening in right-wing media than in centrist, center-left, and left-wing media.

We will make the argument throughout this book that the behavior of the right-wing media ecosystem represents a radicalization of roughly a third of the American media system. We use the term “radicalization” advisedly in two senses. First, to speak of “polarization” is to assume symmetry. No fact emerges more clearly from our analysis of how four million political stories were linked, tweeted, and shared over a three-year period than that there is no symmetry in the architecture and dynamics of communications within the right-wing media ecosystem and outside of it. Second, throughout this period we have observed repeated public humiliation and vicious disinformation campaigns mounted by the leading sites in this sphere against individuals who were the core pillars of Republican identity a mere decade earlier. At the beginning of this period, Jeb Bush, the son and brother of the two most recent Republican presidents, was besmirched as having “close Nazi ties” on Infowars. By November 2017 life-long Republicans who had been appointed to leading law enforcement positions by President George W. Bush found themselves under sustained, weeks-long disinformation campaigns aimed to impugn their integrity and undermine their professional independence. When a solidly conservative party is taken over by its most extreme wing in a campaign that includes attacks that are no less vicious when aimed at that conservative party’s mainstream pillars than they are at the opposition party, we think “radicalization” is an objectively appropriate term.

This radicalization was driven by a group of extreme sites including Breitbart, Infowars, Truthfeed, Zero Hedge, and the Gateway Pundit, none of which claim to follow the norms or processes of professional journalistic objectivity. As we will see time and again, both in our overall analysis of the architecture and in our detailed case studies, even core right-wing sites that do claim to follow journalistic norms, Fox News and the Daily Caller, do not in fact do so, and therefore fail to act as a truth-telling brake on these radical sites. Indeed, repeatedly we found Fox News accrediting and amplifying the excesses of the radical sites. As the case studies in Chapter 5 document, over the course of 2017 Fox News had become the propaganda arm of the White House in all but name. This pattern is not mirrored on the left wing. First, while we do find fringe sites on the left that mirror the radical sites, these simply do not have the kind of visibility and prominence on the left as they do on the right. Second, the most visible sites on the left, like Huffington Post, are at their worst mirrors of Fox News, not of the Gateway Pundit or Zero Hedge. And third, all these sites on the left are tightly integrated

with traditional mainstream media sites like the *New York Times* and the *Washington Post*, and most, though not all, of these sites operate either directly under long-standing journalistic norms or are indirectly sensitive to criticism based on reporting that adheres to such norms. As we show in Chapter 3, there is ample supply of and demand for false hyperpartisan narratives on the left. The difference is that the audience and hyperpartisan commercial clickbait fabricators oriented toward the left form part of a single media ecosystem with center, center-left, and left-wing sites that are committed to journalistic truth-seeking norms. Those norm-constrained sites, both mainstream and net-native, serve as a consistent check on dissemination and validation of the most extreme stories when they do emerge on the left, and have no parallels in the levels of visibility or trust that can perform the same function on the right.

We do not expect our findings to persuade anyone who is already committed to the right-wing media ecosystem. The maps we draw in Chapter 2 could be interpreted differently. They could be viewed as a media system overwhelmed by liberal bias and opposed only by a tightly-clustered set of right-wing sites courageously telling the truth in the teeth of what Sean Hannity calls the “corrupt, lying media,” rather than our interpretation of a radicalized right set apart from a media system anchored in century-old norms of professional journalism. We take up this issue in Chapter 3 where we compare left and right news sites for their patterns of reporting and correction and where we describe our explicit efforts to find conspiracy theories that made it out of the margins of the left to the center of mainstream media. We dedicate Chapter 6 to exploring the modes of failure of mainstream media in their election coverage, and examine the recipients of the Trump Fake News Awards and how they responded to having made the significant errors that won them that honor. We think that fundamentally, anyone who insists on claiming that we cannot draw conclusions about which side is biased, and which side gravitates more closely to the truth, must explain how the media sources most trusted by consistently conservative survey respondents—Fox News, Hannity, Rush Limbaugh, and Glenn Beck—are the equivalent of the sites that occupy the same positions among consistently liberal respondents: NPR, PBS, the BBC, and the *New York Times*.²⁹

The central role of the radicalized right in creating the current crisis of disinformation and misinformation creates a significant challenge for policy recommendations and is not easy to reconcile with democratic theory. It seems too partisan a perspective to convert into a general, nonpartisan policy recommendation or neutral argument about what democracy requires.

And yet, we believe that there is a core set of concerns that transcend party affiliation and should appeal across party lines. First, having a segment of the population that is systematically disengaged from objective journalism and the ability to tell truth from partisan fiction is dangerous to any country. It creates fertile ground for propaganda. Second, it makes actual governance difficult. Other than their major success with tax reform, Republicans found it difficult to govern during the first year of the Trump presidency, despite holding majorities in both houses of Congress and the presidency. In large part, this is due to the inability to bridge the gap between the state of the world as many in their base know it and the state of the world as it is. Third, the divorce of a party base from the institutions and norms that provide a reality check on our leaders is a political disaster waiting to happen—see for instance the primary victory of Roy Moore over Lucas Strange and Moore’s subsequent defeat in the general election. However strident and loyal the party base may be, not even a clear majority of Republican voters is exclusively focused on the right-wing media ecosystem. Over time, the incongruence between the reality inside and outside that ecosystem will make it harder for non-base “lean-Republican” voters to swallow candidates that are palatable inside it. Our hope, then, is that perhaps Republicans see in our findings reason enough to look for a change in the dynamic of the media ecosystem that their most loyal supporters inhabit.

Finding nonpartisan or bipartisan solutions in a society as highly polarized as the United States has become difficult, to say the least. But ignoring the stark partisan asymmetry at the root of our present epistemic crisis will make it impossible to develop solutions that address the actual causes of that crisis. Any argument that depends for its own sense of neutrality and objectivity on drawing empirically false equivalents between Fox News and CNN, much less between top left-wing sites like Mother Jones and Salon and equivalently prominent sites on the right like the Gateway Pundit or Infowars, undermines clear thinking on the problem at hand.

Mainstream Media—Most Americans do not get their news from Facebook, and even most Trump voters did not get their news solely from Fox News or right-wing online media.³⁰ As Thomas Patterson’s study of mainstream media coverage of the 2016 election documented,³¹ and Watts and Rothschild’s study of the *New York Times* in the run-up to the election showed,³² mainstream media coverage of the election was mostly focused on the horserace, was overwhelmingly negative about both candidates, and treated both candidates as equally unfit for office. One of the starkest findings of our work was the

extent to which non-horserace coverage in mainstream media followed the agenda of the right-wing media and the Trump campaign.

As Figure 1.1 shows, Trump got more coverage, and, however negative, the stories still covered his core substantive issues—immigration, jobs, and trade. By contrast, Clinton’s coverage was dominated by scandals—emails, the Clinton Foundation, and, to a lesser extent, Benghazi. The effect of this was most clearly shown in a September 2016 Gallup poll that showed that the word Americans most consistently associated with Clinton was “emails,” followed by “lie,” “scandal,” and “foundation.” By contrast, for Trump, “Mexico,” “immigration,” and so forth were more prominent.³³ Given that only a portion of Trump supporters were primarily focused on the right wing, it would have been impossible for this stark divergence of association to arise without the adoption by major media of the framing and agenda-setting efforts of the right wing and the Trump campaign. Patterson’s explanation for the negativity in coverage is that journalists have developed professional norms that use cynicism and negative coverage as a sign of even-handed, hard-hitting journalism. Certainly the need to attract viewers and earn advertising revenues is rendered easier by focusing on easy-to-digest exciting news like horserace or scandals than by running in-depth analyses of wonkish policy details. None of these are technological or political drivers. They are the basic




FIGURE 1.1 Sentences in mass-market open web media mentioning topics related to Donald Trump and Hillary Clinton during the 2016 election period.

drivers of advertising-supported media. We document and analyze these dynamics in Chapter 6, in particular the fragility of basic journalistic norms of neutrality in the teeth of an asymmetric, propaganda-rich media ecosystem. It is, we believe, impossible to gauge the effects of any of the other actors or dynamics without considering how they flowed through, and affected coverage by, professional mainstream media that were the primary source of news for most Americans outside the right-wing bubble.

Donald Trump: Candidate and President—All the explanations we have presented to this point ignore a central player in the dynamics of the moment—Donald Trump himself. But as we see repeatedly throughout this book, the president played a central catalyzing role in all these dynamics. In Chapter 4, when we look at the topic of immigration and the peaks and valleys in coverage, there is little doubt that Trump launched his campaign with an anti-immigrant message and continued to shape the patterns of coverage with his own statements throughout. Even Breitbart, clearly the most effective media actor promoting immigration as the core agenda of the election and framing it in terms of anti-Muslim fears, seems to have taken its cues from the candidate no less than Trump took cues from the site. Trump's comments as candidate repeatedly drew heated coverage and commentary from the mainstream press. He used this tactic to hold the spotlight from the beginning to the end of the campaign. Audacity, outrage, and divisiveness fueled his campaign. Trump launched his political career largely on his support of the “birther” movement, and has since embraced a wide range of conspiracy theories, from implying that Ted Cruz's father was associated with the Kennedy assassination, through reviving the Vince Foster conspiracy, to asserting that Hillary Clinton aided ISIS. Since becoming president Trump has repeatedly embraced and propagated conspiracy theories against political opponents, many of which fall within the broad frame of the “deep state,” as when he embraced the Uranium One conspiracy theory designed to discredit the FBI and special counsel Robert Mueller's investigation. He has also adopted the term “fake news” to describe all critical mainstream media. It is highly likely his steady drumbeat of condemnation will have only heightened the distrust of media among his supporters.

Trump's use of Twitter has been one the defining facets of his presidency. One way to understand Trump's use of Twitter is as a mechanism to communicate directly with the public, particularly his tens of millions of followers (although it is unknown how many of these are in fact U.S. voters, and how many in fact read his tweets). The picture is more complex than that.

Trump not only uses Twitter to communicate with his followers, but also uses Twitter as a means of exerting power—over the media, the executive branch, the legislature, or opponents. Both as candidate and as president Trump has used Twitter in a feedback loop with the media. There have been several stories that document Trump “live tweeting” cable news, picking up talking points from what the president sees every morning on Fox and Friends, and tweeting them back out to the world.³⁴ There are many signs that guests that appear on the president’s favorite Fox News, for example House Speaker Paul Ryan promoting a spending bill, are in effect performing for an audience of one;³⁵ others have documented how coverage is crafted to influence Trump.³⁶ The influence flows in both directions in an unusual multimedia relationship. Not a day goes by without the president’s tweets becoming a news story. However outlandish a claim, as soon as the president makes or repeats it on Twitter, it is legitimate news not only on right-wing media but across the media ecosystem. In this manner, Trump has been able to insert himself as the center of media attention at any time and with little effort, and in doing so influence the media agenda. During the campaign, this attention “earned” him vastly greater media coverage than Hillary Clinton got. Over the course of his first year in office, the president has used Twitter to short-circuit normal processes in the executive branch as well, such as firing Secretary of State Rex Tillerson, announcing a ban on transgender troops, and creating “facts on the ground” that his subordinates have had to either explain away or, more often than not, accept and adjust their marching orders. While in principle the president could have used regular appearances at the daily press briefing in a similarly mercurial fashion, there is little doubt that the easy availability of a mass media outlet in the palm of his hand, at all hours of the day, has given the president a new and highly unorthodox lever of power.

Overall, the debate since the election seems to have focused on the presidency of Donald Trump more as the consequence to be explained than as an active player in a positive feedback contributing to information disorder. It might also be argued that the exact opposite is the case. That the media dynamics we observe in 2016 and since the election are the anomalous result of the presence of Trump, a charismatic reality TV personality with unusually strong media skills, first as candidate and then as president. There is little doubt that Trump is an outlier in this sense. Nonetheless, the highly asymmetric architecture of the media ecosystem precedes him, as do the asymmetric patterns of political polarization, and we think it more likely that his success was enabled by a political and media landscape ripe for takeover rather than that he himself upended the ecosystem. Trump, as both candidate

and president, was both contributing cause and outcome, operating on the playing field of an already radicalized, asymmetric media ecosystem. As we explain in Chapters 10 and 11, Donald Trump represents the present state of a dynamic system that has been moving Republican politicians, voters, audiences, and media to the right at least since Rush Limbaugh launched this model of mass media propaganda on talk radio in 1988 and became, as the *National Review* wrote in 1993, “the Leader of the Opposition.” In that ecosystem, Trump now operates as catalyst in chief.

Mapping the Actors: Politics, Commerce, Technology, and Centralization

In addition to identifying the specific actors responsible for the current crisis, we try to examine over the course of this book the larger drivers behind those specific actors. Our argument is that the crisis is more institutional than technological, more focused on U.S. media ecosystem dynamics than on Russia, and more driven by asymmetric political polarization than by commercial advertising systems. To highlight these larger questions, we have included two figures below. In both figures, we map each of the actors we have identified above along a horizontal axis of political versus commercial orientation. In Figure 1.2, we map the actors with a vertical axis that distinguishes between threats that come from centralized sources, like states or big companies, as opposed to decentralized sources, like grassroots mobilization or small businesses out to make a buck. In Figure 1.3, we map actors along a vertical axis that distinguishes between threats that are seen as caused by technological change versus those seen as coming from institutional dynamics—laws or social norms that shape how we develop our beliefs about what’s going on and why.

Three things become quite clear from looking at these maps. First, most efforts to understand the apparent epistemic crisis have focused on technology—social media, the Facebook algorithm, behavioral microtargeting, bots, and online organic fragmentation and polarization. Second, Russia and Facebook play a very large role in the explanations. And third, most of the explanations focus on threats and failures that are politically neutral, or nonpartisan.

Our research suggests that our present epistemic crisis has an inescapably partisan shape. The patterns of mistrust in media and lack of ability to tell truth from fiction are not symmetric across the partisan divide. And the

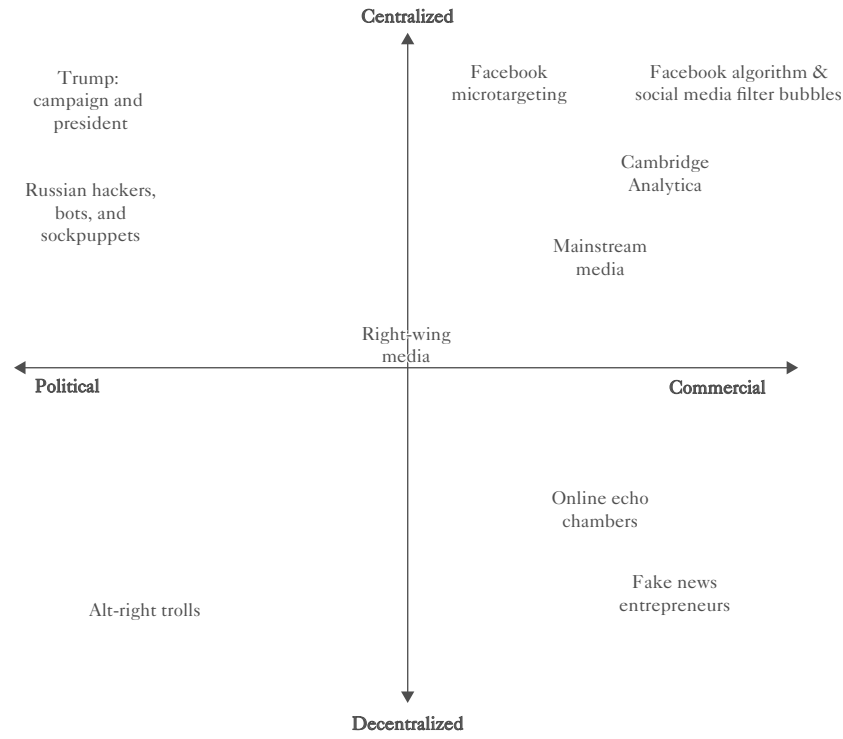


FIGURE 1.2 Threat models distinguished by political vs. commercial orientation and centralized vs. decentralized origins.

fundamental explanation for these differences cannot be laid at the feet of Facebook, Russia, or new technology. They are rooted in long-term changes in American politics. We are not arguing that technology does not matter, that the Russians did not actively interfere with U.S. politics, or that Facebook’s algorithm is unimportant. Instead, we suggest that each of these “usual suspects” acts through and depends on the asymmetric partisan ecosystem that has developed over the past four decades. What that means in practice, for Americans, is that solutions that focus purely on short-term causes, like the Facebook algorithm, are unlikely to significantly improve our public discourse.

For others around the world it means understanding the costs and benefits of proposed interventions based on local institutional conditions, rather than context-free explanations based on “the nature” of the technology. Germany’s new law, which puts Facebook and Google under severe pressure to censor their users to avoid large fines, may be useful for hate speech specifically, but

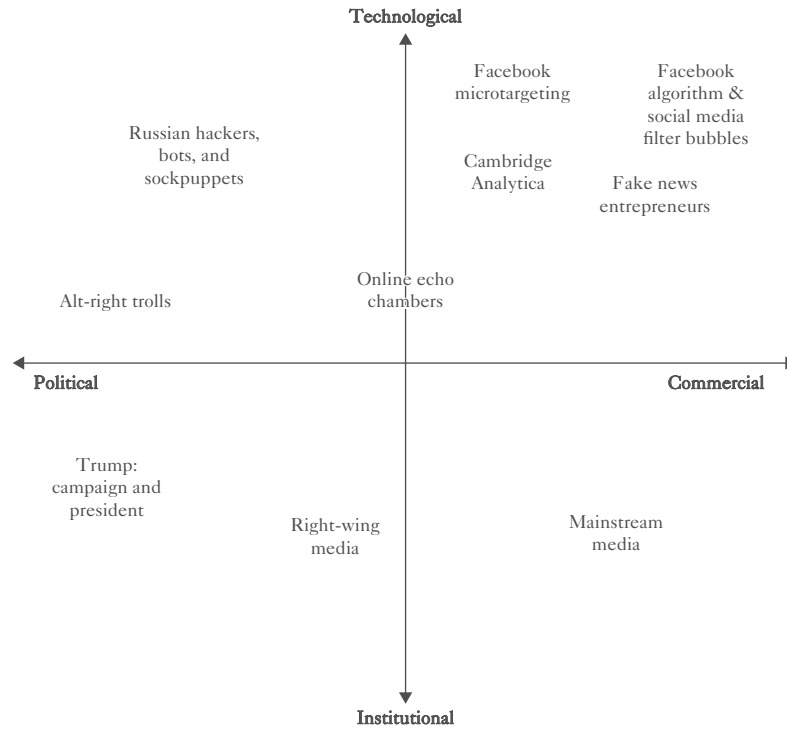


FIGURE 1.3 Threat models distinguished by political vs. commercial orientation and technological vs. institutional origins.

is not likely needed to counter propaganda generally because of the high trust in public media in Germany. This social background fact likely helped keep Russian propaganda on the margins in the 2017 German election, and the highly stable political institutional structures further limited the destabilizing effect of network propaganda and kept the far right at the margins of political significance despite electoral success. On the other hand, perhaps it is precisely Germany's willingness to wade into regulating its public sphere, born of its own bitter historical experience with democratic failure, that has given it its relatively stable media ecosystem. We do not suggest an answer to that question in this book, but we do offer a set of tools and an approach to understanding technology that may help those evaluating that question reach a clearer answer.

Technology does not determine outcomes or the patterns of its own adoption. Specific technologies, under specific institutional and cultural conditions, can certainly contribute to epistemic crisis. Radio during the

Rwandan genocide, and possibly troll-farms in Russia and WhatsApp in India, suggest that technology, in interaction with particular political-institutional conditions, can become the critical ingredient that tips some societies into instability, maybe at local levels or for brief periods. But we have not seen sufficient evidence to support the proposition that social media, or the internet, or technology in itself can be a sufficient cause for democratic destabilization at the national scale. Indeed, our own detailed study of the American case suggests that it is only where the underlying institutional and political-cultural fabric is frayed that technology can exacerbate existing problems and dynamics to the point of crisis. In the 2016 election, it was the already-present asymmetric architecture of news media, and asymmetric attitudes toward professional journalism governed by norms of objectivity, that fed the divergent patterns of media adoption online. It was this asymmetric architecture, and the insularity of right-wing media and audiences from the corrective function of professional journalism, that made that segment of the overall media ecosystem more susceptible to political clickbait fabricators, Russian propaganda, and extremist hyperpartisan bullshit of the Pizzagate variety.

*Definitions: Propaganda and Its Elements,
Purposes, and Outcomes*

The widespread sense that we have entered a “post-truth” era and the general confusion over how we have gotten to this point has led to several careful efforts to define the terms of reference in the debate. The initial surge in “fake news” usage by observers from the center and left was quickly superseded by the term’s adoption by President Trump to denote coverage critical of him, and has since essentially lost any real meaning. Influential work by Clair Wardle and Hossein Derakhshan for the Council of Europe and by Caroline Jack at the Data & Society Institute began to bring order to well-known but ill-defined terms like “*propaganda*,” “*misinformation*,” or “*disinformation*,” as well as introducing neologisms like “*malinformation*” to denote leaks and harassment strategies.³⁷

In the rest of this section, we present a brief history of the study of propaganda and lay out the definitions of these terms as we will use them through the rest of the book. We focus on the information and communications measures, rather than the harassment and intimidation activities. We anchor our definitions both in the salient forms we observed in

our study of communications from April of 2015 to March of 2018 and in the long tradition of propaganda studies.

This segment may be a touch academic for some readers who want to get to the meat of our observations about how propaganda in fact played out in America in the presidential elections and the first year of the Trump presidency. For those readers we offer a brief cheat sheet here and invite you to skip the history and definitions and either go straight to our description of the plan of the book in this chapter or just dive in to Chapter 2.

- “*Propaganda*” and “*disinformation*”: manipulating and misleading people intentionally to achieve political ends.
- “*Network propaganda*”: the ways in which the architecture of a media ecosystem makes it more or less susceptible to disseminating these kinds of manipulations and lies.
- “*Bullshit*”: communications of outlets that don’t care whether their statements are true or false, and usually not what their political effect is, as long as they make a buck.
- “*Misinformation*”: publishing wrong information without meaning to be wrong or having a political purpose in communicating false information.
- “*Disorientation*”: a condition that some propaganda seeks to induce, in which the target population simply loses the ability to tell truth from falsehood or where to go for help in distinguishing between the two.

There are more details in the definitions and history, and like all cheat sheets, this one is neither complete nor precise. But these will serve to make sure that if you skip the next segment you will not miss any important aspects of the chapters that follow.

A brief intellectual history of propaganda

Histories of propaganda, including Harold Laswell’s field-defining work in the 1920s and 1930s,³⁸ emphasize sometimes the ancient origins of the use of communications to exercise control over populations and sometimes the Catholic Congregation for the Propagation of the Faith in the seventeenth century. It seems, nonetheless, that the intense interest in propaganda is a distinctly modern phenomenon. Several factors lent new urgency to the question of how governing elites were to manage mass populations: The discovery of “the masses,” uprooted by industrialization and mass migration as a new object of concern; the emergence of new mass communications

technologies, from the penny presses and mass circulation papers through movies in the nickelodeons to radio in the 1920s; the invention of psychology as a field of scientific inquiry and its application to mass populations; and the urgent need to mobilize these populations in the teeth of total war on a scale never seen before. Propaganda as a field was an application of the modernist commitment to expertise and scientific management, applied to the problem of managing a mass population in time of crisis. Walter Lippmann's words in *Public Opinion* might as well have been written in 2017 about behavioral psychology, A/B testing, and microtargeting as it was in 1922:

That the manufacture of consent is capable of great refinements no one, I think, denies. The process by which public opinions arise is certainly no less intricate than it has appeared in these pages, and the opportunities for manipulation open to anyone who understands the process are plain enough. The creation of consent is not a new art. It is a very old one which was supposed to have died out with the appearance of democracy. But it has not died out. It has, in fact, improved enormously in technic, because it is now based on analysis rather than on rule of thumb. And so, as a result of psychological research, coupled with the modern means of communication, the practice of democracy has turned a corner. A revolution is taking place, infinitely more significant than any shifting of economic power.

Within the life of the generation now in control of affairs, persuasion has become a self-conscious art and a regular organ of popular government. None of us begins to understand the consequences, but it is no daring prophecy to say that the knowledge of how to create consent will alter every political calculation and modify every political premise. Under the impact of propaganda, not necessarily in the sinister meaning of the word alone, the old constants of our thinking have become variables. It is no longer possible, for example, to believe in the original dogma of democracy; that the knowledge needed for the management of human affairs comes up spontaneously from the human heart. Where we act on that theory we expose ourselves to self-deception, and to forms of persuasion that we cannot verify.³⁹

The Committee on Public Information that operated to shape public opinion in support of the American war effort in World War I, the Creel Committee, implemented the idea of applying the best cutting-edge techniques of technology and psychology to engage in the "Engineering of Consent," as one

of its most influential alumni, Edward Bernays, would later call it.⁴⁰ Bernays, in this regard, embodies the translation of cutting-edge social psychology and the idea of expert engineering into the domain of manufacturing public opinion, and became one of the founders of the public relations industry in the early 1920s. Over the course of the 1920s to 1940s, writing about propaganda that remains influential today was caught between the will to systematize the definition and understanding of propaganda so as to make it an appropriately professional, scientific field of practice, and the negative inflection that the term had received during World War I, when all antiwar efforts were branded “German propaganda.” While the negative connotation lingered, the professional orientation toward a managerial definition anchored in psychological manipulation is captured both in Laswell’s classic 1927 definition “Propaganda is the management of collective attitudes by the manipulation of significant symbols”⁴¹ and in the state-of-the-art definition adopted by the Institute for Propaganda Analysis in 1937: “Propaganda is the expression of opinions or actions carried out deliberately by individuals or groups with a view to influencing the opinions or actions of other individuals or groups for predetermined ends and through psychological manipulations.”

The critical elements of this era of professionalized attention to propaganda were, therefore: (a) an actor with intent to manage a (b) target population’s attitudes or behaviors (c) through symbolic manipulation informed by a psychological model of belief or attitude formation and revision, as opposed to rational or deliberative approach. This purposive, managerial approach remains central to self-conscious professional propagandists to this day. The *Army Field Manual of Psychological Operations*, for example, describes the role of PSYOP soldiers to “[i]nfluence foreign populations by expressing information subjectively to influence attitudes and behavior, and to obtain compliance, noninterference, or other desired behavioral changes.”⁴² These definitions emphasize *the propagandist actor*: an agent who acts intentionally; the purpose: to *influence* or manage a *target population*, which distinguishes propaganda from one-on-one persuasion, manipulation, or fraud; and the means: “*manipulation*,” or “expressing information subjectively” in the terms of the PSYOP field manual—that is to say, *communicating* in a manner *behaviorally designed* to trigger a response in the target population to affect *beliefs, attitudes, or preferences* of the target population in order to *obtain behavior compliant* with the subjective goals of the propagandist.

The tension between this understanding of propaganda and a more deliberative or participatory view of democracy was already explicitly present in Lippmann’s 1922 *Public Opinion*. If, as an empirical fact of the matter, the

opinions of citizens as a population are poorly formed and weakly held, and if they are subject to manipulation through ever-more-refined interventions informed by ever-improving scientifically tested social and cognitive psychology, then the idea of deliberative democracy by an informed citizenry exercising self-governance is a utopia. In 1922 Lippmann was still willing to make the argument explicitly, and use it as a basis for an expertise-informed, elite-governed, democracy, recognizing that the inevitability of public opinion being manipulated may as well be used to mobilize support for good policies rather than for exploitation. The same is true for propaganda in markets. Bernays, who had cut his teeth in the Committee on Political Information, would go on to developing marketing campaigns, such as branding cigarettes “torches of freedom” in a 1929 effort to market cigarettes to women.⁴³ If consumer preferences were manufactured by sellers, and citizens’ beliefs were manufactured by elites, then both anchors of liberal market democracies were fundamentally unstable. The use of the term “propaganda” made both of these tensions too palpable, and the term receded from use by expert practitioners, to be replaced by less morally freighted terms: “marketing,” “communications,” “public relations,” or “publicity.” “Propaganda,” when used by mainstream authors, was left to describe what the Soviet Union did. Mostly, its use shifted to become a critical framework from which to criticize modern liberal market society, most famously in Jacques Ellul’s *Propaganda: The Formation of Men’s Attitudes* and later Edward Herman and Noam Chomsky’s *Manufacturing Consent: The Political Economy of the Mass Media*.

Ellul’s now-classic work reoriented the study of propaganda from understanding the practice of intentional management of beliefs and attitudes at the population level to understanding the structure of consciousness in technologically mediated market society. Propaganda was no longer something an actor perpetrates on a population (although it is that too), but the overall social practice and effect that normalizes and regularizes life and meaning in modern, technologically mediated society. “In the midst of increasing mechanization and technological organization, propaganda is simply the means used to prevent these things from being felt as too oppressive and to persuade man to submit with good grace.”⁴⁴ This focus on how propaganda in this pacification sense is a pervasive characteristic in modern, mass-mediated society, and bridging from political communication to marketing and even education, became the typical approach of a mostly dwindling field of study.⁴⁵ Most prominently in this period, Herman and Chomsky’s “Propaganda Model” of mass media offered the most explicit application of the term “propaganda” to the claim that “the media serve, and

propagandize on behalf of, the powerful societal interests that control and finance them.”⁴⁶ They delivered a detailed critique of media coverage of a range of politically sensitive topics and described the dynamics of newsroom politics, of ownership and control, media concentration and advertising dependence. Herman and Chomsky assimilate the term “propaganda” into the more general critique of commercial mass media that became prominent from the 1980s to the early 2000s. Other important contributions in this vein, which did not use the terminology of “propaganda” to make their point, include Ben Bagdikian’s *Media Monopoly*, Neil Postman’s *Amusing Ourselves to Death*, Robert McChesney’s *Rich Media, Poor Democracy*, and Ed Baker’s *Media, Markets, and Democracy*. This literature examined important failings of the commercial mass media model that typified the state of the media pre-internet, and continues to shed light on some of the failures of media conglomerates and the threats of concentrated commercial media to a well-functioning democratic public sphere. But in appropriating the term “propaganda” to describe the broad structure of ideology in modern, technologically mediated market society, as Ellul did, or the failings of commercial mass media during the rise of neoliberalism, as Herman and Chomsky did, the critical turn removed the term “propaganda” from the toolkit of those who wish to study intentional manipulation of public opinion, particularly as applied to politics.

We do not ignore or reject the validity and value of sustained study and critique of the democratic distortions introduced by commercial mass media, particularly concentrated media. Indeed, one of us relied on it in exploring the role of the internet in reversing some of these destructive characteristics of the commercial mass-mediated public sphere.⁴⁷ But the dynamics of the 2016 U.S. presidential election, Brexit, and other political arenas lead us to believe that it would be more useful to adopt the approach of work in the 1990s and 2000s that itself sought to revive the technocratic or scientific study of propaganda as a coherent topic of analysis. This work focused primarily on retaining the negative connotation of propaganda, while overcoming the urge to treat communications from “our” side as “communications,” and those of opponents as “propaganda.” The primary effort was to coherently distinguish “propaganda” from a variety of other terms that refer to communication to a population that has a similar desired outcome: persuasion, marketing, public relations, and education. The most influential treatment in this technical, observational line of work was Gareth Jowett and Victoria O’Donnell’s *Propaganda and Persuasion*.⁴⁸ The book offered the most comprehensive intellectual history of propaganda studies, psychological research on the

influence of media on attitudes and behaviors, and psychological warfare, and sought to systematize a definition and analysis approach out of that broad survey of the field. We follow that line of work in characterizing propaganda. Our claim for these definitions is not their abstract truth but their utility. We think they can help focus analysis on the set of threats that appear to be most salient in the present state and help define a coherent set of research questions without spilling over into too many associated fields of inquiry.

Definitions: Propaganda and Its Elements

Propaganda

Communication designed to manipulate a target population by affecting its beliefs, attitudes, or preferences in order to obtain behavior compliant with political goals of the propagandist.

This definition focuses the study of propaganda on *intentional* communications that are designed by the propagator to obtain outcomes. Unintentionally false communications that in fact affect behavior, but not by design, would not count. It limits the study to communications *targeted at a population*, excluding interpersonal manipulation or very small-scale efforts to manipulate a small group. This limitation in part is to keep the term in its broad original sense of dealing with mass populations, and in part to avoid confounding it simply with all forms of interpersonal manipulation. It limits the term to purposive behavior intended to affect a *political* outcome desired by the propagandist. In this, we purposefully exclude marketing, which fits all other elements of the definition and which shares many of the practices. We do so because we believe normative judgments about how it is appropriate for parties to treat each other in markets may differ from those we consider appropriate for interactions among citizens and between them and their governments or other political actors. Historically, even in more critical usage, “propaganda” has tended to highlight politics and power rather than commercial advantage, although critical usage has tended to expand the meaning and reach of the political. We recognize that what commentators will consider “political” will vary and that a campaign to market soap may be properly described as political if it is understood to reinforce a certain viewpoint about the appropriate standards by which women should be judged, for example. We certainly do not intend to exclude politics in this sense from our definition of political. Nonetheless, we think that confounding political manipulation with manipulative and misleading marketing for commercial

gain will muddy the waters in both. The rise in consumer surveillance and behavioral marketing will certainly require extensive work on measuring and combating such practices. We think it will likely benefit from being part of a renewed literature on misleading advertising and consumer protection and that the study of manipulative marketing and propaganda will share a good bit of empirical overlap in terms of how to identify the practices, measure their effects, and so forth. But we believe there is an advantage to keeping separate the domain of politics, with its normative commitment to democracy, from the domain of commerce, and its normative commitment to welfare, consumer sovereignty, and consumer protection. And using “propaganda” to refer to the political domain is both consistent with longtime usage of the term and allows a more distinct focus for this kind of manipulation from this set of actors who operate within a distinct institutional framework for a set of distinct motivations.

Manipulation

Directly influencing someone’s beliefs, attitudes, or preferences in ways that fall short of what an empathetic observer would deem normatively appropriate in context.

We emphasize the *manipulative* character of communications we deem propaganda. But “manipulation” is itself a term requiring definition. By Cass Sunstein’s account, “manipulation” entails influencing people in a manner that “does not sufficiently engage or appeal to their capacity for reflection and deliberation.”⁴⁹ We think this is too restrictive, because there can be emotional appeals that circumvent the rational capacities but are entirely appropriate for the relevant situation—such as a coach rallying her players’ spirits at halftime after a disastrous first half. The critical issue is appropriateness for the situation. We adopt a variant of Anne Barnhill’s definition that “[m]anipulation is directly influencing someone’s beliefs, desires, or emotions such that she falls short of ideals for belief, desire, or emotion in ways typically not in her self-interest or likely not in her self-interest in the present context.”⁵⁰ We adopt a variant, rather than Barnhill’s definition, because in the political context appeals that are not in the self-interest of the individual are common and appropriate for the context, and so we focus instead on the property of appropriateness for the context. Indeed, we consider “autonomy” rather than “self-interest” as the touchstone, and adopt an “empathetic observer” standard, which one of us initially proposed in the context of characterizing modes of communication as appropriate and inappropriate from the perspective of an autonomous

subject.⁵¹ The “empathetic observer” differs from the “reasonable person” in that she takes the first-person perspective of the target of the communication, and asks whether that person, knowing the entire situation, including the intentions and psychological techniques of the alleged manipulator, would welcome the allegedly manipulative communication. While manipulation is often aided by cutting-edge behavioral psychology, we do not take the fact of scientifically informed design to be definitional, in the sense that manipulation informed by intuition and pop psychology is no less and no more objectionable for that reason.

What “manipulate” adds to a definition of propaganda, which already focuses on intentional action to shape beliefs, attitudes, or preferences, is the need to explain why the communication falls short of a normative ideal for how beliefs, attitudes, or preferences ought to be shaped. Outright false or materially misleading communications are relatively easy to categorize as normatively inappropriate, but emotionally evocative language presents harder questions. What is the difference between Martin Luther King Jr.’s soaring oratory in the “I Have a Dream” speech, and the Gateway Pundit’s anti-immigration article entitled “Obama Changes Law: Allows Immigrants with Blistering STDs and Leprosy into U.S.”? Perhaps appeals to strongly negative emotions such as fear, hatred, or disgust should be considered more inappropriate than strongly positive emotions, such as love, patriotism, or pride? One might imagine that from the perspective of democratic engagement, at least, the former would be more destructive to the possibility of public reason-giving and persuasion than the latter, and so from the perspective of a democracy committed to collective self-governance by people who treat each other as worthy of equal concern as citizens, the latter would more readily fall into the category of normatively inappropriate. But some of the worst abuses in human history were framed in terms that in the abstract sound positive and uplifting, be it love of country and patriotism or the universal solidarity of workers. The empathetic observer allows us to ground the difference in arguments about respecting the autonomy of members of the population subject to the intervention. Its risk, of course, is that those who think it is trivially true to think that King’s oratory is appropriate and the Gateway Pundit’s is not have to contend with the possibility that the readers of the Gateway Pundit are fully aware of the intent and effect of the communication, and desire it no less than the nearly defeated athletes at halftime desire the rousing pep talk from the coach. If that turns out to be sociologically true, then one needs some framework based not on respect for the autonomy of citizens as individuals but based on a more collective normative framework, such as what democracy requires of citizens.

We do not try to resolve that question here but emphasize that some form of manipulation is a necessary part of justifying the normatively negative connotation of “propaganda” and that connotation must have a well-defined normative foundation other than “I don’t agree with what they said.”

Disinformation

Dissemination of explicitly false or misleading information.

We use the term “disinformation” to denote a subset of propaganda that includes dissemination of explicitly false or misleading information. The falsehood may be the origin of the information, as when Russian-controlled Facebook or Twitter accounts masquerade as American, or it may be in relation to specific facts, as when Alex Jones of Infowars ran a Pizzagate story that he was later forced to retract. We mean to include both “black” and “gray” propaganda in the term disinformation, that is to say, propaganda whose source or content is purely false, as well as propaganda whose source and content is more subtly masked and manipulated to appear other than what it is.

Bullshit—Here we adopt philosopher Harry Frankfurt’s by now widely popular definition, which covers communicating with no regard to the truth or falsehood of the statements made. A liar knows the truth and speaks what he knows to be untruthful. The bullshit artist “does not care whether the things he says describe reality correctly. He just picks them out, or makes them up, to suit his purpose.”⁵² “Fake news” producers in the original meaning, that is, purely commercial actors with no apparent political agenda who propagated made-up stories to garner engagements and advertising revenue, were the quintessence of bullshit artists in this sense. Not all bullshit is propaganda, because its propagators are indifferent to its impact on beliefs, attitudes, or preferences and have no agenda to shape behavior, political or otherwise, other than to induce a click. Here, in our case studies, we repeatedly encounter behavior that may be pure bullshit artistry and may be propaganda, depending on our best understanding of the motive and beliefs of the propagator. Broadcasting propaganda can be enormously profitable. It is hard to tell whether Sean Hannity or Alex Jones, for example, were acting as propagandists or bullshit artists in some of the stories we tell in Chapters 4 through 9. In all these cases, and those of many other actors in our studies, there is certainly a symbiotic relationship between the propagandists and the bullshit artists, with each providing the other with bits of narrative and

framing that could be used, copied, circulated, and amplified throughout the network, such that as a practical matter what we interchangeably call commercial bullshit and political clickbait fabricators formed an important component of network propaganda.

Network Propaganda—What we observe in our broad, macroscale studies as well as in our detailed case studies is that the overall effect on beliefs and attitudes emerges from the interaction among a diverse and often broad set of discrete sources and narrative bits. The effects we define below—induced misperceptions, disorientation, and distraction—which contribute to population-scale changes in attitudes and beliefs, come not from a single story or source but from the fact that a wide range of outlets, some controlled by the propagandist, most not, repeat various versions of the propagandist’s communications, adding credibility and improving recall of the false, misleading, or otherwise manipulative narrative in the target population, and disseminating that narrative more widely in that population. We call this dynamic “network propaganda” to allow us to treat it as a whole, while, as network analysis usually does, emphasizing the possibility and necessity of zooming in and out from individual nodes and connections to meso- and macroscale dynamics in order to understand the propagation characteristics and effect of propaganda campaigns, and in particular the role of network architecture and information flow dynamics in supporting and accelerating propagation, as opposed to resisting or correcting the propagandist efforts as they begin to propagate.

Propaganda Feedback Loop—This is a network dynamic in which media outlets, political elites, activists, and publics form and break connections based on the contents of statements, and that progressively lowers the costs of telling lies that are consistent with a shared political narrative and increases the costs of resisting that shared narrative in the name of truth. A network caught within such a feedback loop makes it difficult for a media outlet or politician to adopt a consistently truth-focused strategy without being expelled from the network and losing influence in the relevant segment of the public. We dedicate Chapter 3 to describing and documenting the effect.

Propaganda Pipeline and the Attention Backbone—In earlier work, one of us identified a dynamic, “the attention backbone,” whereby peripheral nodes in the network garner attention for their agenda and frame, attract attention to it within a subnetwork of like-minded users who are intensely

interested in the perspective, and through relatively more visible members of that subnetwork amplify that agenda or frame progressively upward to ever-more publicly influential sites to reach the public at large.⁵³ He argued that this dynamic was democratizing in that it circumvented the power of major media outlets to set, or sell the power to set, the agenda, and that it allowed for more diverse viewpoints than would normally be admissible in the mass-mediated public sphere to rise to the public agenda. In our work here, we have encountered that same structural dynamic being used to propagate disinformation and propaganda from peripheral nodes into the core. While the mechanisms are identical, we distinguish its use for propaganda purposes (which is determined by the character of the content, not its transmission mechanism) as “the propaganda pipeline” and offer an analysis and case study in Chapter 7.

Effects of Propaganda

Induced Misperceptions

Politically active factual beliefs that are false, contradict the best available evidence in the public domain, or represent patently implausible interpretations of observed evidence.

By “misperception,” we follow D.J. Flynn, Brendan Nyhan, and Jason Reifler, who define “misperceptions” as “factual beliefs that are false or contradict the best available evidence in the public domain.”⁵⁴ By “politically active,” we follow Jennifer Hochschild and Katherine Levine Einstein’s typology, and focus on what people think they know that is associated with distinctive involvement in the public arena.⁵⁵ This limitation flows from our concern in propaganda with politics. Our observations regarding the practices of disinformation in the American public sphere lead us to emphasize that the falsehood of a belief cannot be limited to factual misstatement but must include patently implausible interpretations of observed evidence. We include as “patently implausible” both interpretations containing obvious logical errors and, more controversially, bad faith interpretations or framing that contradicts well-documented professional consensus. While interpretation adds fuzziness to judgments about beliefs, it is impossible to ignore the broad range of disinformation tactics we have observed that construct paranoid conspiracy interpretations around a core of true facts. Limiting our focus to discrete, false statements of fact will cause us to lose sight of more important and more invidious manipulations than simply communication of wrong

Jerry Brown Signs Bill Allowing Illegal Immigrants to Vote - Breitbart
www.breitbart.com/california/.../gov-jerry-brown-signs-bill-allowing-illegal-aliens-vot... ▼
 Oct 12, 2015 - On Saturday, California Governor **Jerry Brown** signed Assembly Bill 1461, the New Motor Voter **Act**, which **will** automatically register people to **vote** through the DMV, and could result in **illegal aliens voting**. Any person who renewed or secured a driver's license through the DMV may now register to **vote**, ...

FIGURE 1.4 Google search result for Breitbart story on California motor voter law.

facts. For example, in one of its most widely shared immigration-related stories, Breitbart wrote a factually correct set of statements, that Governor Jerry Brown of California signed a motor voter law that, if the secretary of state fails to check the citizenship of applicants, could allow undocumented immigrants to vote. The story was posted with the URL designation: “Jerry Brown Signs Bill Allowing Illegal Immigrants to Vote.” This then shaped how the headline appeared in Google searches (Figure 1.4) or the Facebook newsfeed, so that this was the title as the story appeared in users’ normal flow, creating what we consider an example of a false interpretation of a set of true facts. While the text of the article included factually correct statements—that Jerry Brown signed a law that automatically registered eligible voters when they obtained a driver’s license, and that, if the secretary of state of California should fail to check the eligibility of registrants such a law could result in undocumented immigrants being registered—the gap between the chain of unlikely but possible events that would lead to such an outcome and the highly salient interpretation embedded in the distribution of the story suggests an effort of intentional manipulation that would lead to “misperception” as we define it here.

Distraction—Propaganda can often take the form of distracting a population into paying no attention to a given subject and thus losing the capacity to form politically active beliefs about or attitudes toward it. As Gary King, Jennifer Pan, and Margaret Roberts showed in a large empirical study of Chinese domestic propaganda that the “50 Cent Army”—the tens or hundreds of thousands of people paid by the Chinese government to post online to affect online discussion—primarily engages in distraction, rather than persuasion or manipulation to achieve particular results. Studying over 400 million posts by these paid propagandists, King, Pan, and Roberts argue that they do not generally engage anyone in debate, or attempt to persuade anyone of anything; instead, their primary purpose appears to be to “distract and redirect public attention from discussion or events with collective action potential.”⁵⁶

Disorientation—The emphasis on disorientation appears in the literature on modern Russian propaganda, both in inward-focused applications and in its international propaganda outlets, Sputnik and RT (formerly, Russia Today). Here, the purpose is not to convince the audience of any particular truth but instead to make it impossible for people in the society subject to the propagandist’s intervention to tell truth from non-truth. As Peter Pomerantsev put it in *Nothing Is True and Everything is Possible*:

In today’s Russia, by contrast, the idea of truth is irrelevant. On Russian “news” broadcasts, the borders between fact and fiction have become utterly blurred. Russian current-affairs programs feature apparent actors posing as refugees from eastern Ukraine, crying for the cameras about invented threats from imagined fascist gangs. [...] “The public likes how our main TV channels present material, the tone of our programs,” [the deputy minister of communications] said. “The share of viewers for news programs on Russian TV has doubled over the last two months.” The Kremlin tells its stories well, having mastered the mixture of authoritarianism and entertainment culture. The notion of “journalism,” in the sense of reporting “facts” or “truth,” has been wiped out. [...] When the Kremlin and its affiliated media outlets spat out outlandish stories about the downing of Malaysia Airlines Flight 17 over eastern Ukraine in July—reports that characterized the crash as everything from an assault by Ukrainian fighter jets following U.S. instructions, to an attempted NATO attack on Putin’s private jet—they were trying not so much to convince viewers of any one version of events, but rather to leave them confused, paranoid, and passive—living in a Kremlin-controlled virtual reality that can no longer be mediated or debated by any appeal to “truth.”⁵⁷

As we will see, disorientation has been a central strategy of right-wing media since the early days of Rush Limbaugh’s emergence as a popular conservative radio talk show host and political commentator with millions of listeners. Limbaugh’s decades-long diatribes against one or all of what he calls “the four corners of Deceit”—government, academia, science, and the media—seem designed to disorient his audience and unmoor them from the core institutionalized mechanisms for defining truth in modernity. Repeatedly throughout our research for this book we have encountered truly fantastical stories circulated widely in the right-wing media ecosystem, from Hillary Clinton trafficking in Haitian children to satisfy her husband’s unnatural

lusts, to Hillary Clinton herself participating in pedophilia on “Orgy Island,” to John Podesta’s participation in satanic rituals, to the Uranium One story in which the special counsel investigating Russian interference in support of Donald Trump’s campaign, Robert Mueller, and the deputy attorney general who appointed him, Rod Rosenstein, were portrayed as corruptly facilitating the Obama administration’s sale of 20 percent of America’s nuclear capabilities to Russia. These are all stories reported widely in the core sites of the right wing, and polls report that substantial numbers of Republicans claim to believe these stories—whether because they actually believe them factually or because claiming to believe them is part of what identifies them as Republicans.⁵⁸ But all of these seem so ludicrously implausible that it is difficult to imagine that they are in fact intended to make people believe them, rather than simply to create a profound disorientation and disconnect from any sense that there is anyone who actually “knows the truth.” Left with nothing but this anomic disorientation, audiences can no longer tell truth from fiction, even if they want to. They are left with nothing but to choose statements that are ideologically congenial or mark them as members of the tribe. And in a world in which there is no truth, the most entertaining conspiracy theory will often win.

Propaganda and its effects, in our usage, always involve intention to communicate. A potentially important fact about the internet is that it may also create a good bit of error and confusion that are not the result of anyone’s effort to shape political belief. We follow other work in defining misinformation.

Misinformation

Communication of false information without intent to deceive, manipulate, or otherwise obtain an outcome.

This may occur because in the 24-hour news cycle, journalists and others who seek to keep abreast of breaking news make honest mistakes; it may occur because the internet allows a much wider range of people to communicate broadly, and it covers many people without the resources or training to avoid error or correct it quickly. A major line of concern with the internet and social media is precisely that what characterizes our “post-truth” moment is misinformation in this sense. A common articulation of this concern is that as a society we are experiencing the disorientation that results from inhabiting a fast-moving, chaotic process and that the information chaos is an emergent property of a ubiquitously connected, always-on society over which we have no control.

Our data lead us to believe that, at least insofar as political communications are concerned, this diagnosis is false. The much simpler explanation in the political domain is that we are operating in a propaganda-rich environment and that network propaganda is a much deeper threat to democracy than any out-of-human-control emergent socio-technical process. We recognize that misinformation may play a larger role in fields that are not pervasively populated by intentional, well-organized, and well-resourced actors. False beliefs about health, such as anti-vaccine and anti-fluoridation sentiment, may well be more of function of misinformation than propaganda dynamics, as Brittney Seymour and her collaborators have shown.⁵⁹ But political communication operates in its own distinct ecosystem and that network is pervaded by propaganda. Describing the evidence that leads us to this conclusion and documenting the highly asymmetric pattern of susceptibility to and diffusion of propaganda and bullshit in the American media ecosystem will take up most of the book. As we will see, our conclusions make it difficult to identify solutions that are consistent with free speech and respect for freewheeling political contestation. We identify some responses, incremental and partial though they may be. We do believe that there are discrete interventions that can help. But the fundamental challenge is not purely or even primarily technological. It is institutional and cultural; which is to say, ultimately, political.

Plan of the Book

In the next chapter we describe our macrolevel findings about the architecture of the American news media ecosystem from 2015 to 2018. We collected and analyzed two million stories published during the 2016 presidential election campaign, and another 1.9 million stories about the Trump presidency during its first year. We analyze patterns of interlinking between the sites to understand the relations of authority and credibility among publishers high and low, and the tweeting and Facebook sharing practices of users to understand attention patterns to these media. What we find is a highly asymmetric media ecosystem, with a tightly integrated, insular right wing, while the rest of the media ecosystem comprises not a symmetrically separated left and a center but rather a single media ecosystem, spanning the range from relatively conservative publications like the *Wall Street Journal* to liberal or left publications like *The Nation* or *Mother Jones*, anchored in the traditional professional media. This basic pattern has remain unchanged in the year since Donald Trump took office, and, as we show in Chapter 11, was likely

already in place at least as early as the 2012 election cycle. Chapter 3 presents a model of how such an insular media ecosystem might emerge and how two fundamentally different media ecosystems can coexist—one in which false narratives that reinforce partisan identity not only flourish but crowd out true narratives even when these are presented by leading insiders, and the other in which false narratives are tested, confronted, and contained by diverse outlets and actors operating in a truth-oriented norms dynamic. We then show how parallel but politically divergent false rumors, about Trump raping a 13-year-old and Hillary and Bill Clinton being involved in pedophilia, followed fundamentally different paths through the media ecosystems into which each was introduced. We argue that the difference in how the two highly divergent media ecosystems amplified or resisted the false narratives, not in the initial availability of falsehoods or the enthusiasm with which audiences wanted to hear them, made the Trump rape allegations wither on the vine while the Clinton pedophilia rumors thrived.

In Part Two we analyze the main actors in the asymmetric media ecosystem and how they used it to affect the formation of beliefs and the propagation of disinformation in the American public sphere. Chapter 4 looks at how the more radical parts of the right-wing media ecosystem, centered on Breitbart, interacted with Donald Trump as a candidate to force immigration to the forefront of the substantive agenda in both the Republican primary and the general election, and framed immigration primarily in terms of Islamophobia and threats to personal security. Chapter 5 turns to Fox News in particular. Through a series of case studies surrounding the central controversy of the first year of the Trump presidency—the Trump Russia investigation—we show how Fox News repeatedly mounted propaganda attacks at major transition moments in the controversy—the Michael Flynn firing in March 2017, when Fox adopted the “deep state” framing of the entire controversy; the James Comey firing and the Robert Mueller appointment in May 2017, when Fox propagated the Seth Rich murder conspiracy; and in October and November 2017, when the arrests of Paul Manafort and guilty plea of Michael Flynn seemed to mark a new level of threat to the president, when Fox reframed the Uranium One story as an attack on the integrity of the FBI and Justice Department officials in charge of the investigation. In each case, the attacks persisted over one or more months. In the case of the “deep state,” the attack created the basic frame through which the right-wing media ecosystem interpreted the entire investigation. In each case the narrative and framing elements were repeated across diverse programs, by diverse personalities, across Fox News and Fox Business on television and on Fox’s online site.

Chapter 6 turns its attention to the failure and recovery modes of mainstream media. Most Americans do not get their news from the right-wing media ecosystem. The 2016 election was heavily influenced by the ways in which those parts of the public sphere anchored around mainstream media operated. In particular, we show how standard journalistic norms of balance and chasing the scoop and the flashy headline create predictable failure modes in the presence of a highly asymmetric and propaganda-rich ecosystem. In the search for balance, professional media outlets emphasized negative coverage and focused heavily on scandals in their coverage of Clinton, particularly on emails and the Clinton Foundation. In search of the scoop, traditional media could not resist email-based coverage, and repeatedly covered it with a mismatch between the actual revelations, which were usually fairly banal, and the framing and headlines, which often overstated the significance of the findings. We use the case of the Clinton Foundation coverage in particular as a detailed case study of how Steve Bannon, Breitbart-CEO turned Trump campaign CEO and later White House chief strategist for the first seven months of the Trump administration, manipulated first the *New York Times*, and then most of the rest of the mainstream media ecosystem, into lending their credibility to his opposition research efforts in ways that were likely enormously damaging to Clinton's standing with the broad electorate. We end Chapter 6 with a study of the several news failures that populated the president's Fake News Awards list for 2017 and show that in all these cases the media ecosystem dynamic of "the rest," outside the right wing, functioned to introduce corrections early and to reinforce the incentives of all media outlets and reporters in that larger part of the ecosystem to commit resources and follow procedures that will allow them to get their facts straight.

Part Three looks at the remaining prime suspects in manipulating the election or political coverage during the election and the first year of the Trump presidency. In Chapter 7 we document a widely reported, but we think relatively less important part of the story—the propaganda pipeline. We show how statements by marginal actors on Reddit and 4chan were collated and prepared for propagation by more visible sites; and we show how this technique was used by both alt-right and Russia-related actors to successfully get a story from the periphery to Sean Hannity. But we also explain why we think that this process depends on actors to which we dedicated Chapters 3, 4, and 5, and why the propaganda pipeline, while open, was of secondary importance.

Chapter 8 addresses Russian propaganda more generally. We describe the existing sources of evidence that overwhelmingly support the proposition

that Russia mounted sustained and significant information operations in the United States. We do not raise doubts as to the fact that the efforts occurred and continue. We do, however outline a series of questions about how important the Russian interference really was. Repeatedly, when we assess the concrete allegations made, we suggest reasons to think that these were likely of marginal significance, and mostly jumped on a bandwagon already hurtling down whatever path it was taking rather than giving that bandwagon the shove that got it going. Understanding not only that Russian propaganda efforts happened, but also how effective they were is, we think, critical to putting those efforts in context. Critically, if the biggest win for Russian information operations is to disorient American political communications, in the way we define the term here, then overstating the impact of those efforts actually helps consolidate their success. We do not suggest that we should understate the effects when we do see them. But it is important not to confuse the high degree to which Russian operations are observable with the extent to which they actually made a difference to politically active beliefs, attitudes, and behaviors on America.

Chapter 9 looks at the three major Facebook-based culprits—fake news entrepreneurs, behavioral manipulation by Cambridge Analytica, and plain old Facebook microtargeted advertising. We review the work of others, and complement it to some extent with our own data, to explain why we are skeptical of the importance of clickbait fabricators or Cambridge Analytica. The core long-term concern, we suggest, is Facebook’s capacity to run highly targeted campaigns, using behavioral science and large-scale experimentation. In Chapter 13 we address possible solutions regarding political advertising regulation and the adoption, by regulation or voluntary compliance, of a public health approach to providing independent research accountability to the effects of Facebook and other major platforms on the health of the American media ecosystem.

Part Four takes on the claim that “the internet polarizes” or that social media creates “filter bubbles” as the primary causal mechanism for polarization online, as compared to the claim that the highly asymmetric media ecosystem we observe is a function of long-term institutional and political dynamics. Chapter 10 engages with the political science literature on polarization and shows that polarization long precedes the internet and results primarily from asymmetric political-elite-driven dynamics. We believe that that elite process, however, is influenced by the propaganda feedback loop we describe in Chapter 3. We lay out that argument in Chapter 11, when we turn to media history and describe the rise of second-wave right-wing media. We take a

political economy approach that explains how institutions, politics, culture, and technology combine to explain why Rush Limbaugh, televangelism, and Fox News were able to emerge as mass media when they did, rather than remaining, as first-generation right-wing media after World War II had, small niche players. We then describe how the emergence of the online right-wing media ecosystem simply followed the offline media ecosystem architecture and, indeed, was left little choice by the propaganda feedback loop but to follow the path that it took. We see that asymmetric polarization precedes the emergence of the internet, and that even today the internet is highly unlikely to be the main cause of polarization, by comparison to Fox News and talk radio. In Chapter 12 we turn to consider what these insights imply for how we think about the internet; whether it can, or cannot, contribute to democratization, and under what conditions. Chapter 13 offers solutions. Unfortunately, the complex, long-term causes we outline do not lend themselves to small technocratic solutions. But we do emphasize adaptations that traditional media can undertake, in particular shifting the performance of objectivity from demonstrating neutrality to institutionalized accountability in truth seeking, as well as reforms in rules surrounding political advertising and data collection and use in behavioral advertising. We consider these all to be meaningful incremental steps to at least contain the extension of the unhealthy propaganda dynamics we observe to deeper and more invidious forms, but we acknowledge that given the decades-long effects of the propaganda feedback loop on the architecture of right-wing media in America, the solution is only likely to come with sustained political change.

In the final chapter, we underscore two core conclusions. First, if we are to understand how technology impacts society in general, and politics and democratic communications in particular, we must not be caught up in the particular, novel, technical disruption. Instead, we have to expand our viewpoint across time and systems, and understand the long term structural interactions between technology, institutions, and culture. Through this broader and longer-term lens, the present epistemic crisis is not made of technology; it cannot be placed at the feet of the internet, social media, or artificial intelligence. It is a phenomenon rooted in the radicalization of the right wing of American politics and a thirty-year process of media markets rewarding right-wing propagandists. We suspect that a similarly broad and long-term lens will be required to properly understand the rise of far-right parties and their information ecosystems elsewhere. At least in the United States, we find here that a failure to do so results in severe misdiagnosis of the challenges we face.

Second, much of contemporary discussion of the causes of crisis confound what is novel and observable with what actually has impact. It is hard to find bots; and when researchers find them, they are new and interesting. But repeatedly in our work they operate as background noise, and do not change the structure of the conversation. It is hard to spot Russians; and when Russia-hunters find them, they are sinister and newly menacing. But our work suggests that, while they were there, and much of what they did is likely illegal, they were mostly jumping on bandwagons already hurtling full-tilt downhill and cheering on a process entirely made in America. It is challenging to measure, and titillating to imagine, young nihilists generating fake news for American consumption and overwhelming traditional media. But when we measure the actual impact, it does not seem to be significant. So too with psychographically-informed behavioral marketing Cambridge Analytica-style.

It is not complacency that we seek to communicate, but the necessity of evidence-based diagnosis. If the evidence is too narrowly focused, or lacks the context to interpret the observations correctly, it will lead us to misdiagnose the problem and to develop solutions for emotionally salient but functionally marginal contributing causes of information disorder. These will, in turn, divert our attention and efforts to the wrong solutions, for the wrong reasons, at the wrong time. In the present crisis of the project of democracy, that is a misdiagnosis that we cannot afford.