Digital Influence Warfare
in the Age of Social Media

James J. F. Forest

PRAEGER SECURITY INTERNATIONAL

Digital Influence Warfare in the Age of Social Media

Digital Influence Warfare in the Age of Social Media

JAMES J. F. FOREST

Praeger Security International

PRAEGER®

An Imprint of ABC-CLIO, LLC
Santa Barbara, California • Denver, Colorado

Library of Congress Cataloging-in-Publication Data

Names: Forest, James J. F., author.
Title: Digital influence warfare in the age of social media / James J.F. Forest.
Description: Santa Barbara, California : Praeger, An Imprint of ABC-CLIO, LLC, [2021] | Series: Praeger security international | Includes bibliographical references and index.
Identifiers: LCCN 2020054801 (print) | LCCN 2020054802 (ebook) | ISBN 9781440870095 (hardcover) | ISBN 9781440870101 (ebook)
Subjects: LCSH: Information warfare. | Social media—Political aspects. | Disinformation. | Mass media and propaganda. | Mass media and public opinion.
Classification: LCC UB275 .F67 2021 (print) | LCC UB275 (ebook) | DDC 355.3/437—dc23
LC record available at https://lccn.loc.gov/2020054801
LC ebook record available at https://lccn.loc.gov/2020054802

ISBN: 978-1-4408-7009-5 (print)
978-1-4408-7010-1 (ebook)

25 24 23 22 21 1 2 3 4 5

This book is also available as an eBook.

Praeger
An Imprint of ABC-CLIO, LLC

ABC-CLIO, LLC
147 Castilian Drive
Santa Barbara, California 93117
www.abc-clio.com

This book is printed on acid-free paper ♾

Manufactured in the United States of America

Contents

Preface *vii*

Acknowledgments *xiii*

1. An Introduction to Digital Influence Warfare 1
2. Goals and Strategies: Influencing with Purpose 29
3. Tactics and Tools: Technical Dimensions of Digital Influence 67
4. Psychologies of Persuasion: Human Dimensions of Digital Influence 111
5. Exploiting the Digital Influence Silos in America 153
6. Information Dominance and Attention Dominance 189
7. Concluding Thoughts and Concerns for the Future 221

Notes 245

Index 297

Preface

Let me begin with an honest self-reflection. I have published more than 20 books, and this has been among the most difficult of all, in part because of the tumultuous events swirling around us during the time this enters the publisher's review and production process (e.g., the COVID-19 pandemic and related economic turmoil, the nationwide protests against police brutality, the ongoing threat of foreign and domestic terrorism, and a highly polarizing presidential election). This has also been an unusually difficult topic to write about because of the emotions it provokes, such as dismay, frustration, anger, powerlessness, and even hopelessness—all in response to the fact that we have been (and continue to be) attacked on a daily basis by malicious actors, both foreign and domestic, who want to use our online information access and sharing activities as weapons against us.

The research and writing of this book required an extensive journey of discovery, and when I began the journey in late 2017, one of my goals was to find some answers to a puzzling question. I had recently seen an increasing number of people I know—people whom I consider reasonable and intelligent—expressing opinions and beliefs that I knew to be untrue, things that could not be supported by any factual evidence. This was occurring sometimes in face-to-face conversations, but much more so in discussions online, and particularly on social media. Why would these people be so convinced of something that is proven completely false by all factual evidence? Further, when factual evidence was presented to them clearly proving that they were incorrect, these people would just turn away and continue repeating their support of the falsehoods to anyone

who would listen. Or, in several instances, they would try to argue that their beliefs were more valid than the facts.

What was going on? These were not stupid people, and they did not exhibit the signs of someone who had been brainwashed (whatever that word really means) by a cult or terrorist group. Yet they had come to embrace particular narratives about a range of issues and people that the rest of the world rejected. Having studied terrorism and terrorists for nearly 20 years, I thought I had a fairly good handle on things like extremism and radicalization. One of my books—*Influence Warfare: How Terrorists and Governments Fight to Shape Perceptions in a War of Ideas* (Praeger, 2009)—had even examined various aspects of propaganda, psychological operations, and disinformation, with particular focus on how websites, blogs, email, online videos, digital magazines, and other such things were used to shape beliefs, attitudes, and behaviors. My primary research question at that time was how governments were competing with terrorists for influence and support in the public domain, and particularly on the Internet. But a decade later, I have now come to realize that the scope of this earlier research was far too limited: what we see today is a much broader and complex terrain, in which the rapid advancement and global usage of social media has introduced new concepts, strategies, and tactics for influence warfare that did not exist just a decade ago, and a much broader range of actors are using these strategies and tactics than ever before.

So, for the past few years I have been studying this broader phenomenon of what I now call digital influence warfare—reading an ever-growing stack of books, academic research articles, reports by government agencies and think tanks, and much more. Many of these publications have focused on Russia's massive disinformation efforts aimed at the populations of countries like Estonia, Ukraine, the United Kingdom, and the United States. But an increasing number of other countries are also engaged in similar activities, including China, Iran, Saudi Arabia, and Turkey. As discussed in the introductory chapter, one report from the Oxford Internet Institute found that in 2018 there were disinformation efforts of one kind or another in 70 countries around the world. But at the same time, extremists and terrorists have also taken advantage of new opportunities for provoking fear—even livestreaming videos of attacks in Kenya and New Zealand—with dramatic results. And a profit-generating business model has shifted the entire landscape of influence warfare in a new—and decidedly more dangerous—direction, especially during the worldwide COVID-19 pandemic of 2020. In today's attention economy, the ability to shape perceptions and influence behavior through social media is a major source of power and profit.

After finding many examples of state and non-state actors using the many new tools of digital influence, I also began to appreciate the strategic mechanics of it—the psychology of persuasion or social influence

applied in online environments. From my previous work, I could understand several of the relevant concepts already, like exploiting a person's "fear of missing out" (FOMO), extremist ideologies, dehumanization, in-group indoctrination, out-group "othering," provocation, propaganda, psychological operations, political warfare, website defacement, and tools for manipulating photos and videos. I knew that repetition and framing were important elements of effective communication, and I appreciated the dangers of conspiracy theories. I also knew something about data mining, social network analysis, and algorithms, having co-taught a course at West Point on information warfare many years ago. However, there were other terms that I was learning for the first time, like trolling, doxxing, gaslighting, hashtag flooding, deepfakes, astroturfing, ragebait, digital information silo, and so forth.

What I found in my research were many studies that said basically the same thing: there are clear and recognizable strategies and tactics being used by certain people to manipulate the perceptions and behaviors of others. Generally speaking, three kinds of people are involved: influencers, enablers, and targets. Some of the publications I encountered used terms like "influence aggressor" to describe the individuals whose actions are described in this book. They may be state sponsored or driven by ideological beliefs, profits, and many other kinds of motives. Their ability to identify advantageous targets for influence efforts has become easier based on all the information that is available about us. As we'll examine in several chapters of this book, billions of people worldwide are providing free and unfiltered access to themselves by posting photos, personal revelations, telling people where they are at a given moment, and showcasing who their friends and family are. Further, because of the profit models that pervade the attention economy, Internet firms track a user's identity and patterns of behavior so they can formulate the right kinds of advertising campaigns. Just as every click and keystroke can be monitored, recorded, and used for analysis that generates advertising profits for the Internet companies, the same data can inform a digital influence strategy.

The targets of digital influence efforts have become increasingly accessible as well, particularly those who engage more frequently on social media and other online information resources on a daily basis. Influencers can now use Facebook or other social media platforms to pinpoint much more precisely the types of individuals who might be receptive to the information (or misinformation) they want to disseminate. The targets could be virtually anyone, but influencers quickly find that they'll have more success by choosing targets whose beliefs and values indicate certain predispositions and biases. Further, changing a target's mind about something may be an objective, but this is much more difficult than finding targets whom you only need to nudge a little in a certain direction

or simply confirm for them that their biases and prejudices about others are justified. Effective influencers have learned how to capitalize on the fact that the Internet provides the means to shape a reality that caters to the disposition of its users. And while the target is often described as an unwitting participant (or even victim) in digital influence warfare, this not always the case. As we'll see reflected in several chapters of this book, many individuals are actively seeking out disinformation and fake sources of information online solely for the purpose of providing confirmation for what they want to believe.

For their part, the digital influencer could pursue any number of goals and objectives. Some may want to deceive, disinform, and provoke emotional responses (including outrage) in order to influence certain people's voting behavior. Other influencers may want to strengthen the commitment of the target's beliefs, reinforcing their certainty and conviction in something; this may include attacking scientific evidence that supports inconvenient truths. The goals of digital influence also include various forms of online recruitment efforts by global jihadists and other terrorist networks, as well as the nurturing of online communities whose members embrace extremist ideologies (e.g., white nationalism, neo-Nazism, sovereign citizens, ANTIFA, or the incel movement). Sometimes, the strategy can involve convincing the targets that what they believe or think they know is based on false information. Or the strategy could be to convince the target that what "other people" believe or think they know is based on false information, leading to a sense of superiority over those naive "others." A particularly powerful form of digital influence involves convincing targets that the beliefs and convictions they are particularly passionate about are severely threatened by other members of society and must be defended. Similarly, a goal of a digital influence effort could be to encourage broader patterns of questioning and uncertainty, leading the targets to believe that nothing is true and anything may be possible. This in turn creates opportunities for the spread of disinformation and conspiracy theories. And other online influencers may simply want to market and sell products, services, and ideas.

There are also a variety of tactics involved in digital influence warfare, from deception (including information deception, identity deception, and engagement deception) to emotional provocation and outright attacking the target (including bullying, hacking, exposing embarrassing information online, etc.). We'll examine these and much more in chapter 3. But across this diversity of goals and tactics, what most of them have in common is that they are intended to shape the perceptions and behaviors of targets in ways that will benefit the influencers more than the targets. In other words, the influencer rarely has the best interests of the target in mind. This seems to hold true regardless of whether the goals of the influencer are political, economic, social, religious, or other categories of belief and behavior.

And finally, in addition to the relative ease of identifying and accessing viable targets, the influencer can also monitor and assess the impact of their influence effort by gathering and analyzing data on the target's reception and reaction to the information they were exposed to. Success in digital influence warfare can be measured by the target's behavior. Did they do something that the influencer wanted them to—for example, vote, buy, protest, join, reject, or some other behavioral response? Did they express some kind of emotional response (outrage, anger, sympathy, encouragement, etc.)? With this assessment in hand, the influencer can then refine their efforts to maximize effectiveness.

With all these developments in mind, I thought a book focused on digital influence warfare would be useful for academics, policymakers, and the general public. The chapters of the book are organized around a series of questions I sought to answer during my intellectual journey through the research on this topic. My search for answers led me through a ton of published research on the psychology of persuasion, in which experts have identified a wide variety of ways in which ordinary individuals can be persuaded—and in some cases, even to do some terrible things to other people. I also revisited the history of influence warfare, with particular focus on Russia and its Active Measures program. Along the way, I found an emerging body of research about what I now call digital influence mercenaries, and I found many examples of non-state actors who are profiting by deceiving and provoking people on social media. A separate book on that topic is now in the works. My journey also led me to the research on technological tools used by state and non-state actors in their digital influence efforts. As a result, I know more now about deepfake images and videos than a person of my technical incompetence should know.

There also seems to be widespread agreement in the published materials on this matter that something ought to be done to curb malicious uses of social media (and other forms of online information and interaction). Social media platforms are certainly doing more today than they were in 2016 to curb the malicious kinds of digital influence efforts described in this book. But I've come to the conclusion that each of us as individual citizens has a responsibility as well. When we stop and think about the influencers behind the information we see and hear, we tend not to be as open to exploitation. Further, these influence attempts—both foreign and domestic—should make us angry: for the most part, there is no informed consent; nobody asked us for our permission to deceive or manipulate us. So, we should get angry enough to do something about it.

We should also expect greater commitment from our government for policies and public education to confront these issues. Digital influence warfare represents a form of cyberattack that requires more than network systems firewall and security. Confronting and deflecting these digital influence efforts require a kind of societal firewall, a psychological barrier

of shared resistance and resilience that rejects and defeats these attempts. Only when a society proves completely invulnerable to digital influence attacks will there be a true deterrent. In the absence of that, our enemies will continue trying.

While this book was being written, our nation endured acrimonious political campaigns, a rising tide of right-wing anti-government extremism, and the deadly COVID-19 virus spreading to countries around the world. Various forms of social mediated disinformation, disorientation, and conspiracy theories have accompanied these and many other major events. Reflecting on this now, it becomes clear to me that unfortunately I chose to research and write a book about a topic where things have been very fast moving and ever-changing. By the time this volume hits the shelves in 2021, some of the analyses and recommendations contained within may be overtaken by events. I ask your indulgence and understanding for this.

As I mentioned at the outset, the research and writing of this book required an extensive journey of discovery, and to be honest, much of what I discovered was rather unpleasant. I have learned more about the darker elements of psychology and human nature—and about technology, social media algorithms, deviant mercenaries, and much more—than I had originally thought possible. I have written and rewritten several chapters multiple times, reorganized the entire volume at least a dozen times, and even scrapped entire chapters (some of which may appear someday as articles or essays in different publications). I have had to go outside my own fields of education, counterterrorism, and international security studies for material used in this book, including such disciplines as psychology, sociology, information technology, criminal justice, communication, political science, and many others. In the course of integrating various information from these disciplines, it was of course necessary to summarize research findings and concepts, so to the experts in those fields who may feel slighted that I overlooked their important contributions, I apologize. In embracing the ethos of the curious mind, I have encountered numerous things about our modern world in recent years that have proved deeply disturbing to me. My academic training prompted me to document these things over the course of several years and eventually (with the prompting of a publisher) put pen to paper in an effort to make sense of it all. Thus, this book represents the product of an intellectual adventure, an account of where I looked for answers and what I learned along the way. I should conclude here with a warning that readers may experience mild whiplash between research-based theories on political, psychological, and influence warfare and my personal observations or whimsical attempts at humor. I hope you enjoy the roller-coaster ride and find the book worthwhile.

Acknowledgments

I owe considerable gratitude to literally thousands of people who have significantly influenced my intellectual journal over the past two decades. Some of them I consider friends and colleagues, while others I have never even met. Some have been coworkers or guests who lectured in my courses, and even coauthored publications with me, while others have only communicated with me briefly online. But they have also helped me answer questions and find new perspectives. The abbreviated list I'd like to especially thank includes: Alex Schmid, Andrew Silke, Annette Idler, Arie Perliger, Assaf Moghadam, Bill Braniff, Brian Fishman, Brian Jenkins, Bruce Hoffman, Colin Clarke, Clint Watts, Daniel Byman, David Kilcullen, David Ronfeldt, Dorothy Denning, Emerson Brooking, Eric Schmitt, Erica Chenoweth, Gabi Weimann, Gary LaFree, GEN Wayne Downing, Greg Miller, Henry Crumpton, J.M. Berger, Jacob Shapiro, Jade Parker, Jarret Brachman, Jennifer Giroux, Jessica Stern, Jim Duggan, Joe Felter, John Arquilla, John Horgan, Joshua Gelzer, Juan Merizalde, Kurt Braddock, Martha Crenshaw, Matthew Levitt, Maura Conway, Max Abrahms, Michael Hayden, BG (ret.) Michael Meese, Michael Sheehan, Nada Barkos, Neil Shortland, Paul Cruikshank, Peter Neumann, Peter W. Singer, Richard Shultz, Robert Cialdini, Rolf Mowatt-Larssen, BG (ret.) Russell Howard, Ryan Evans, Sheldon Zhang, Thom Shanker, Thomas Fingar, Tom Nichols, Walter Lacquer, William McCants, and ADM (ret.) William McRaven.

There are also a growing number of experts and organizations in this emerging field of what I am loosely calling digital influence studies, and I have benefitted enormously from many of them in researching and writing this book. If you are interested in the contents of this book, you will find

the works of these people most enlightening, particularly the hardworking folks at the Oxford Internet Institute's Computational Propaganda Project, Graphika, the Global Network on Extremism and Technology, the Centre for the Analysis of Social Media, the Rand Corporation's Truth Decay project, and the Stanford Internet Observatory. Weekly publications like *First Draft*, *Popular Information*, and *The Source* (published by the Atlantic Council's Digital Forensic Research Lab) are strongly recommended. I also recommend following the online commentary on these and other topics by Barb McQuade, Cass Sunstein, Carl Miller, Caroline Orr, Cindy Otis, Claire Wardle, Emma Barrett, Emma Briant, Erin Gallagher, Jay Rosen, Joan Donavan, Judd Legum, Kate Starbird, Marc Owen Jones, Natalia Antonova, Nathaniel Gleicher, Nick Carmody, Olga Belogolova, Peter Pomerantsev, Phil Howard, Samantha Bradshaw, Yael Eisenstat, and others followed by the Twitter account @DIWbook. And I'm grateful to Naomi Shiffman and her colleagues at CrowdTangle for showing me how to analyze data on the spread of disinformation via social media platforms.

I also want to include here a special shout-out to professors and mentors in my graduate school many years ago who took me under their wing and showed me how I could make potentially worthwhile contributions to the academic profession, especially Patricia Gumport (at Stanford University) and Philip Altbach (at Boston College). I also greatly appreciate my former colleagues at the U.S. Military Academy. I learned so much during my nine years there, particularly from my friends and colleagues in the Department of Social Sciences and the Combating Terrorism Center, as well as from the faculty in the Department of Electrical Engineering, with whom I collaborated on teaching an information warfare course for several years.

I thank the publisher, Praeger/ABC-CLIO, and particularly the editorial staff and proofreaders who helped ensure this was not a complete literary disaster. And finally, I express my appreciation to my family members: Alicia, Chloe, Jack, John, Jason, Jeremy, Jody, Jesse, Jael, and Mary. They are all positive sources of influence in my life, and I am forever grateful.

CHAPTER 1

An Introduction to Digital Influence Warfare

During the process of researching and writing this book, various friends and colleagues would ask me to explain what the term "digital influence warfare" really means and why I chose this term for the book's title. Admittedly, I haven't always had the most articulate way of responding to this question, so let's begin this introductory chapter by providing my best effort to define and explain the term. First, let's consider what each word means:

Digital: Anything online, anything you see on a computer, smartphone, etc. is inherently digital—in other words, composed of digits (1s and 0s) that form text, pixels, sound, etc. We are surrounded by an online ecosystem of digital information providers and tools, from websites, blogs, discussion forums, and targeted email campaigns to social media, video streaming, and virtual reality. While various strategies and tactics of influence warfare have existed for centuries, this book focuses on new and emerging digital forms of it, and the technological environments that enable unique tactical innovations in manipulative behavior.

Influence: An ability to convince others to think or do something. Drawing from decades of research in psychology, marketing, education, sociology, and other disciplines, we have learned the most effective ways an information provider can persuade other people, to shape their beliefs in ways that lead them to embrace one perspective and reject others, and to adopt behaviors in alignment with that perspective. Some influence efforts are intended to strengthen existing beliefs, while others may try to challenge and change those beliefs, but in general the underlying goal of most influence efforts is to impact the behaviors that are driven by what people believe. Influence operations exploit information systems (like social media platforms) to manipulate audiences for the purpose of achieving strategic goals (including political, social, and economic). Sometimes the influencer wants to change people's views and behaviors, while in other instances, they want to strengthen existing beliefs rather than change them (e.g., amplifying existing levels

of distrust and divisions within a society). Throughout the many examples provided in this book, one entity is using information (or in many cases disinformation) in order to gain the power to influence another.

Warfare: A type of human behavior that involves winning and losing, in which there are attackers and targets, offensive and defensive strategies, tactics and weapons. There are also frequently innocent victims, and sometimes third party allies and mercenaries are involved. Warfare is a means to an end, usually some sort of political objectives pursued by the aggressor. It can be a way of gaining power and/or diminishing the power of others—for example, a goal could be to degrade the functional integrity of a democratic society that is considered an adversary or peer competitor.

So, the combination of these three concepts gives us the term "digital influence warfare." In short, it refers to the landscape of online psychological operations, information operations, and political warfare through which a malicious actor (state or non-state) achieves its goals by manipulating the beliefs and behaviors of others. It involves the use of persuasion tactics, information and disinformation, provocation, identity deception, computer network hacking, altered videos and images, cyberbullying, and many other types of activity explored in this book. Examples of digital influence warfare range from using armies of trolls to flood a social media platform with a narrative or view on a specific (often social or political) issue to using thousands of computer-generated accounts ("bots") to manufacture the perception of massive support for (or opposition to) something or someone. And while there has been much attention in the media about Russia and China engaging in these activities, there are both foreign and domestic examples of influence warfare.

The central goal of influence warfare is—and has always been—fairly straightforward: the attacker wants to shape or reshape the reality in which the target believes in order to achieve some sort of strategic objective.[1] However, the context in which this "weaponization of information" takes place has changed significantly over the past two decades. The rise of the Internet and social media companies, whose profit model is based on an "attention economy," has been a game changer. Within the attention economy, the most valued content is that which is most likely to attract attention, with no regard to whether it is beneficial or harmful, true or untrue. New tools have emerged for creating and spreading information (and disinformation) on a global scale. Connectivity in the digital realm is now much easier, and yet—as we'll examine later in this book—ironically the emergence of hyperpartisan influence silos has sequestered many online users into separate communities who reject the credibility and merits of each other's ideas, beliefs, and narratives.

This is why fake information can be so readily believed—as long as it is tailored to support what you want to believe, it will be believed. And it

has never been easier to tailor information of all kinds for a specific audience online. In later chapters of this book, we'll examine the role of deepfake images and videos, memes, fake websites, and many other tools used in digital influence operations. But in this introductory discussion, let's review some important points about terms and terminology and look at a small handful of examples that illustrate what this book is about. Then the latter part of this chapter will provide a brief overview of what readers will find in the rest of the book.

COMPARING INFLUENCE WARFARE WITH INFORMATION OPERATIONS

How does the term digital influence warfare relate to other similar terms like "information operations" or "information warfare"? Indeed, there are reports published on these topics every year, some of which do address the issue of influencing targets. However, those terms have also been increasingly used to describe computer network attacks (often by highly trained military units) like hacking into databases to observe or steal information, pervert information, or replace some kinds of information with other information, and so forth. Traditional military uses of the term "information warfare" have also focused on protecting our own data from those kinds of attacks by adversaries. Of course, computer network attacks like these can certainly be used to send a message (e.g., about a target's vulnerabilities and the attacker's capabilities), and in that way, they could be a means of influencing others. States may want "information dominance" over the populations of other states. This would include computer network operations, deception, public affairs, public diplomacy, perception management, psychological operations, electronic countermeasures, jamming, and defense suppression.[2] Similar terms in this broad landscape include public diplomacy and strategic communications. Cyber operations and cybersecurity have also been intertwined with discussions about information operations.

I prefer to use the term "influence warfare" to describe the kinds of activities in which the focus is not on the information but on the *purposes* of that information, that is, propaganda, misinformation, disinformation, and other kinds of efforts in which the implicit goal of the information is to shape perceptions and influence behavior. Further, influence warfare strategies and tactics—particularly as we have seen online—also involve more than just manipulation of information; they can include behavior signaling (swarming and bandwagoning), trolling, gaslighting, and other means by which the target is provoked into having an emotional response that typically overpowers any rational thought or behavior. Clickbait, memes, and ragebait (for example) are not really seen as a form of information

operations as normally conceived, but it is certainly a means of influencing others via the Internet.

Similarly, other terms addressing the overall concept of cybersecurity can be somewhat confusing. Many of you are familiar with the concept of computer hacking, but this is different. While there is some conceptual overlap, the term "digital influence" warfare should not be confused with terms like "cyberwar," in which the attacker seeks to damage the functionality of information technology, computer systems, and communication networks. Other terms traditionally associated with using computers to attack others including cybersecurity, cyberterrorism, and information warfare—and even digital warfare—are not really what this book is about. Those terms usually apply to attacking other countries' critical infrastructure and military computer systems using tools for hacking into—and degrading or even destroying the functional integrity of—those systems.

But unlike conventional cyberattacks, the goal of a digital influence warfare campaign is not about degrading the functional integrity of a computer system. Rather, it is to use those computer systems against the target in whatever ways might benefit that attacker's objectives. Often, those objectives include a basic "divide and conquer" strategy—a society that is disunited will fight among themselves over lots of things instead of coming together in the face of a threat that only some of them believe is there. The emphasis is thus on the middle word "influence," where a broad diversity of activities are meant to shape the perceptions, choices, and behaviors of a society—and in some cases, the goal may in fact be to make the target dysfunctional as a society. This is not simply propaganda, fake news, or perception manipulation. It is a battle over what people believe is reality and the decisions that each individual makes. The victors in this battle are the attackers who have convinced scores of victims to make decisions that directly benefit the attackers.

As Michael Erbschloe explains, a difference between cyberwarfare and what he describes as "social media warfare" is that "cyber warfare requires a far higher level of technical knowledge and skill. Social media warfare is easier to learn and faster to deploy; but effective social media warfare, like cyber warfare, requires discipline and long-term dedication for successful deployment or defense."[3] Competency in digital influence warfare can be measured by one's ability to successfully influence perceptions and behaviors through information provided by digital means. In 1997, Charles Swett—the Acting Deputy Director for Low Intensity Conflict Policy in the U.S. Department of Defense—offered a warning about how the future would include uses of the Internet "for spreading propaganda by extremist groups and disinformation about U.S. activities."[4] Unfortunately, we have seen much more than those kinds of influence efforts. Nation-states have attempted to impact democratic elections in other countries, as well

as manipulate the perceptions of their own populations. Attempts to provoke outrage and sow discord among a society have increased dramatically, especially with the rise of social media. Today, the Internet offers a unique information environment that brings many advantages to influence warfare campaigns, or what the Atlantic Council's Digital Forensics Lab[5] refers to as "cyber-enabled influence operations." A recent Soufan Center analysis observed how "all modern conflict now features a significant and growing social media component, an extension of the propaganda that has accompanied war for ages."[6]

Meanwhile, the Oxford Internet Institute refers to contemporary forms of influence warfare as "computational propaganda,"[7] while other researchers have examined the rising threat of "information aggressors" and "information wars—sometimes aimed at persuasion, often morphing into vicious cyberbullying."[8] Books, research articles, and reports have been published over the last decade describing the "age of weaponized narrative,"[9] "media manipulation,"[10] or "information disorder."[11] One report portrays the threat as "an increasingly hostile series of aggressive actions between opposing groups . . . [while] wars—though among virtual communities—pit states against states, states against non-state actors, and networks of non-state groups against similar networks."[12] Throughout these digital influence wars, the attacker wins (and the target loses) by successfully influencing the target to think and do things that benefit the attacker's political, social, or other goals. We will examine these and other goals at length in chapter 2, along with a number of prominent examples of digital influence efforts—some of which you have likely seen in your own social media account at some point.

THE "WARFARE" PERSPECTIVE

Using the terminology of warfare when describing information operations and digital influence efforts can be confusing, but there are many precedents to consider. Richard Stengel (a former Undersecretary of State and editor of *Time* magazine) chose the term "information wars" for the title of his recent book,[13] and Danah Boyd (a Principal Researcher at Microsoft and the founder/president of *Data & Society*) used the same term to describe a variety of influence efforts in 2017.[14] This term has also been used at the very top of the Kremlin: Vladimir Putin's spokesperson, Dmitry Peskov, openly says that Russia is in a state of "information war,"[15] and Vyascheslav Volodin, the deputy head of Putin's administration, views social media as a battlefield.[16] In 2014, NATO's top military commander Philip Breedlove called the disinformation campaign around Russia's annexation of Crimea "the most amazing information warfare blitzkrieg we have ever seen in the history of information warfare."[17] A Senate Armed Services Committee hearing in March 2017 addressed "Russian influence

and unconventional warfare operations."[18] And a 2018 UNESCO report describes how "the 21st century has seen the *weaponization of information* on an unprecedented scale. Powerful new technology makes the manipulation and fabrication of content simple, and social networks dramatically amplify falsehoods peddled by States, populist politicians, and dishonest corporate entities, as they are shared by uncritical publics"[19] (my emphasis added).

In a 2019 report by the Rand Corporation, the authors "propose the term *virtual societal warfare* to capture the emerging reality . . . [involving] informational mechanisms of coercion and manipulation."[20] They explain how:

> This warfare involves the use of largely nonkinetic, information-based aggression to attack the social stability of rival nations. It is virtual because, for the most part, these strategies do not employ direct physical violence or destruction. (This concept, therefore, excludes both direct military attack as well as large-scale cyberattacks designed to wreak havoc on a nation's physical infrastructure and cause actual damage.) It is societal because both the targets and the participants in such campaigns stretch across society, and because the goal is to undermine the efficient functioning, levels of trust, and ultimately the very stability of the target society. And it is warfare because, in its potentially more elaborate forms, it represents an activity designed to achieve supremacy over rival nations, not merely to gain relative advantage in an ongoing competition but to gain decisive victory in ways that leave the target nation subject to the attacker's will.[21]

In choosing to use the terminology of warfare for this book, part of my reasoning was the recognition of a particular aggressiveness in the use of social media, and the Internet more generally, to attack targets on behalf of political goals. If war is the continuation of politics by other means—"a real political instrument, a continuation of political commerce," as Carl von Clausewitz suggested[22]—then it would seem appropriate to view the tactics and strategies described here as a form of warfare. Further, war is never an isolated act, but rather is a means to achieve specific goals and objectives over time. Wars require some sort of defensive measures taken by those being targeted, and inevitably, there are casualties of war. Failure to adopt the most effective measures in response to these adversaries could be disastrous for the future of truthful discourse and civil democracy, as Nina Jankowicz explains in her book *How to Lose the Information War*.[23]

Other relevant literature published in recent years also incorporate the language of warfare, including *War in 140 Characters: How Social Media is Reshaping Conflict in the Twenty-First Century*, *Social Media Warfare*, and *LikeWar: The Weaponization of Social Media*.[24] In *LikeWar*, authors Singer and Brooking describe how "the Internet is a battlefield . . . a platform for achieving the goals of whichever actor manipulates it most effectively. Its

weaponization, and the conflicts that then erupt on it, define both what happens on the Internet and what we take away from it. Battle on the Internet is continuous, the battlefield is contiguous, and the information it produces is contagious. The best and worst aspects of human nature duel over what truly matters most online: our attention and engagement."[25] Similarly, other publications have referred to "Cyber troops"[26] engaged in various kinds of activities and the need to fight against these influence efforts in the "digital trenches."[27]

Additional terms closely associated with influence warfare include "political warfare," which was used by the legendary diplomat George Kennan in 1948 to describe "the employment of all the means at a nation's command, short of war, to achieve its national objectives. Such operations are both overt and covert . . ." and can include various kinds of "propaganda" as well as covert operations that provide clandestine support to underground resistance in hostile states.[28] Paul Smith describes political warfare as "the use of political means to compel an opponent to do one's will," and "its chief aspect is the use of words, images, and ideas, commonly known, according to context, as propaganda and psychological warfare."[29] Carnes Lord notes a "tendency to use the terms psychological warfare and political warfare interchangeably" along with "a variety of similar terms—ideological warfare, the war of ideas, political communication and more."[30] And if you are interested in the topic of this book, you will surely enjoy Thomas Rid's excellent book *Active Measures: The Secret History of Disinformation and Political Warfare*.[31]

Altogether, there are political, psychological, and informational dimensions to what Brad Ward refers to as "strategic influence."[32] A recent report by the Rand Corporation explains how "information warfare . . . works in various ways by amplifying, obfuscating, and, at times, persuading" and observes that "political warfare often exploits shared ethnic or religious bonds or other internal seams."[33] Another term frequently encountered in this realm is "strategic communications," which U.S. military reports have defined as "focused efforts to understand and engage key audiences in order to create, strengthen, or preserve conditions favorable for the advancement of interests, policies, and objectives through the use of coordinated programs, plans, themes, messages, and products synchronized with the actions of all instruments of national power."[34] According to a 2004 Defense Science Board report:

> Strategic communication requires a sophisticated method that maps perceptions and influence networks, identifies policy priorities, formulates objectives, focuses on "doable tasks," develops themes and messages, employs relevant channels, leverages new strategic and tactical dynamics, and monitors success. This approach will build on in-depth knowledge of other cultures and factors that motivate human behavior. It will adapt techniques of skillful political campaigning,

even as it avoids slogans, quick fixes, and mind sets of winners and losers. It will search out credible messengers and create message authority . . . It will engage in a respectful dialogue of ideas that begins with listening and assumes decades of sustained effort.[35]

My 2009 book *Influence Warfare* referred frequently to a "strategic communications battlespace" as "the contested terrain upon which all types of information from competing sources seeks to influence our thoughts and actions for or against a particular set of objectives."[36] Terms like "battlespace" and "warfare" may seem odd when the discussion centers on information and influence. However, as we'll see in later chapters of this book, what I found in the course of my research indicates that the weaponization of information—particularly of an emotionally provocative nature—has become a major problem worldwide. In his national best-selling book *Influence: The Psychology of Persuasion*, author Robert Caldini chose to title the very first chapter as "Weapons of Influence." This was no coincidence—the same terminology of weapons or weaponization has been used on countless occasions to describe things used to attacks our thoughts and beliefs. For example, in their book *Age of Propaganda: The Everyday Use and Abuse of Persuasion*, Anthony Pratkanis and Elliot Aronson urge their readers to consider the metaphor of "propaganda is invasion" (i.e., an attacker is trying to conquer a target audience's minds and beliefs).[37]

As a recent European Commission report observes, "Disinformation strategies have evolved from 'hack and dump' cyber-attacks, and randomly sharing conspiracy or made-up stories, into a more complex ecosystem where narratives are used to feed people with emotionally charged true and false information, ready to be 'weaponized' when necessary. Manipulated information . . . [enables] rewriting reality, where the narration of facts (true, partial or false) counts more than the facts themselves."[38] In May 2019, Brian Jenkins—an internationally respected expert in national security—chaired a workshop of Cold War-era subject matter experts on Russian information warfare, veterans who had served in the White House, the State Department, the United States Information Agency, the Pentagon, the FBI, and the intelligence community. The title of his 47-page report from this event? *Russia's Weapons of Mass Deception*.[39] And in a September 2019 *Time* magazine article, former U.S. State Department senior official Richard Stengel argued that "we are all actors in a global information war that is ubiquitous, difficult to comprehend and taking place at the speed of light. . . . Governments, non-state actors and terrorists are creating their own narratives that have nothing to do with reality. These false narratives undermine our democracy and the ability of free people to make intelligent choices."[40]

From this perspective, one might assume the U.S. government has developed some sort of "influence warfare strategy" to defend our nation

from such attempts. To date, that does not exist, although the U.S. Department of Defense has specific definitions for several of the terms associated with influence warfare, including the following:

- *Information Operations (IO)*: "The integrated employment of electronic warfare, computer network operations, psychological operations, military deception, and operations security, in concert with specified supporting and related capabilities, to influence, disrupt, corrupt or usurp adversarial human and automated decision making while protecting our own."[41] Information operations can also help enable a commander to interrupt or stop the flow of information to their adversaries.[42]
- *Psychological Operations (PSYOP)*: Efforts to convey selected truthful information and indicators to foreign audiences to influence their emotions, motives, objective reasoning, and ultimately, the behavior of their governments, organizations, groups, and individuals. The purpose of PSYOP is to induce or reinforce foreign attitudes and behavior favorable to the originator's objectives.[43] PSYOP employs various media such as magazines, radio, newspapers, television, email, dropping leaflets on adversarial territory, and so forth.[44]
- *Information Warfare*: The offensive and defensive use of information and information systems to deny, exploit, corrupt, or destroy an adversary's information, information-based processes, information systems, and computer-based networks while protecting one's own. Such actions are designed to achieve advantages over military, political, or business adversaries.[45]

A recent report by the think tank Demos provides a different working definition of "information operations" as "a non-kinetic, coordinated attempt to inauthentically manipulate an information environment in a systemic/strategic way, using means which are coordinated, covert and inauthentic in order to achieve political or social objectives."[46] The term "propaganda"—described by Jason Stanley as "a means to strengthen and spread the acceptance of an ideology"[47]—is also used frequently to describe influence warfare efforts. According to Philip Zimbardo et al., the main purposes of propaganda have included attempts to weaken or change the emotional, ideological, or behavioral allegiance of individuals to their group (army, unit, village, nation, etc.); to split apart component subgroups of the enemy to reduce their combined effectiveness; to ensure compliance of civilian populations in occupied zones; and to refute an effective theme in the propaganda of the enemy.[48]

Of course, the use of influence operations to achieve the goals of political and psychological warfare is well known among those who study military history, foreign affairs, and international security. As we'll briefly review in chapter 2, influence warfare was prominent in both world wars

and throughout the Cold War. And while Russia continues to invest heavily in its Active Measures program of *Dezinformatsiya*,[49] they are far from being the only states investing in these kinds of activities. Further, it is important to note that this is not only a state-based phenomenon. Influence warfare is also seen in many contemporary insurgencies—in fact, as Thomas Hammes noted, "Insurgent campaigns have shifted from military campaigns supported by information operations to strategic communications campaigns supported by guerilla and terrorist operations."[50] While combating the Malaya insurgency in November 1952, British Field Marshal Sir Gerald Templer observed that "the shooting side of the business is only 25% of the trouble and the other 75% lies in getting the people of this country behind us."[51] This is why modern terrorist groups and extremist movements have also used the spread of propaganda and disinformation for their own purposes, as described most recently in Kurt Braddock's book *Weaponized Words*.[52] Examples range from al-Qaeda's online magazine *Inspire* and the Islamic State's version *Dabiq* to Hezbollah managing its own satellite television network *Al-Manar* and Hamas flooding the Internet with images and videos during their 2012 conflict with Israel.[53] We'll review several specific instances of this in chapter 2.

To sum up, we will explore together throughout this book many different aspects of a modern technology-based form of influence strategies and tactics that have been used by states and non-state actors for centuries. Manipulating the perceptions and behaviors of large populations is made more efficient and effective by a wide range of Internet and social media platforms. It is clear now why Steve Bannon (head of Brietbart at the time) told his staff that the Internet was not just a communications medium; it was a "powerful weapon of war."[54] It is also clear why Facebook recently launched a team responsible for detecting and disrupting the kinds of disinformation campaigns that we have all seen in the news lately, and according to one account, the members of this team view their social media platform as "terrain" where war is waged. Members of this team view themselves as "defenders" against malicious "attackers" whom they must force "downhill" to a position of weakness.[55]

For all these reasons, the term "digital influence warfare" is used in this book to describe a variety of strategies (e.g., to confuse, disorient, destabilize, and increase doubt and uncertainty), tactics (e.g., the promulgation of disinformation and fake news, spoofing, spamming, and hasthtag flooding), and tools (e.g., automated troll farms, "sock-puppet" networks, hijacked accounts, and deepfake images and video). It is a term that draws from the kinds of activities that are often described as political warfare, information warfare, asymmetric warfare, and cyberwarfare as well as strategic communications, information operations, psychological operations, public relations, marketing, and behavioral manipulation. Throughout all of these kinds of activities, we find a common goal: to influence

the thoughts, actions, and reactions of human targets. The fact that an aggressor is seeking to influence the target against their wishes, in order to achieve certain strategic goals, leads us to consider this a form of warfare.

EXAMPLES OF DIGITAL INFLUENCE WARFARE

To illustrate what digital influence warfare is, let's review some specific examples. The next several pages of this chapter will provide just a small handful, and while other chapters of this book contain dozens more, these are all just a small representative sampling of a much larger and diverse landscape of both foreign and domestic influence efforts. Many of the examples provided in this book are linked directly to Russia—and most observers will agree that this country has been the most active and aggressive state sponsor of these types of activities in recent years. In fact, as we'll examine in chapter 2, Russia has a long history of disinformation operations, from the early days of the Soviet Union and throughout the Cold War. According to Mark Galeotti, "In the immediate aftermath of the Crimean seizure, the notion of a radically new style of hybrid war fighting took the West by storm, and led to both insightful analysis and panicked caricatures. This has been called 'new generation warfare,' 'ambiguous warfare,' 'full-spectrum warfare' or even 'non-linear war.'"[56] Galeotti takes issue with the overuse of the term "hybrid warfare," preferring instead to focus on "political warfare" as the overall framework that best describes Russia's engagement with its perceived adversaries. Throughout all the books, government reports, and scholarly journal articles on what Russia does (and why) in pursuit of its foreign policy objectives, the term "warfare" is used quite frequently, and in my view quite logically. The government leaders of Russia are clearly engaged in a war to influence perceptions, beliefs, and behaviors of others. It's a war without bullets or tanks, but instead different kinds of weapons are used—especially weaponized information—and there are clear examples of aggression, targets, defenders, tactics, strategies, goals, winners, losers, and innocent victims. The main thrust of its Active Measures program today takes place online, largely (but not exclusively) via social media platforms, using the tactics and tools of digital influence (described in chapter 3) to sow confusion and spread disinformation about its invasion of Ukraine and annexation of the Crimean peninsula, the shooting down of Flight MH17, and many other issues.

Many of Russia's digital influence campaigns have directly targeted the United States. For example, on June 8, 2016, a Facebook user calling himself Melvin Redick, a family man from Harrisburg, Pennsylvania, posted a link to DCLeaks.com and wrote that users should check out "the hidden truth about Hillary Clinton, George Soros and other leaders of the US." The profile photograph of "Redick" showed a middle-aged man in

a baseball cap alongside his young daughter—but Pennsylvania records showed no evidence of Redick's existence, and the photograph matched an image of an unsuspecting man in Brazil. U.S. intelligence experts later announced, "with high confidence," that DCLeaks was a fake news website created by Russia's military-intelligence agency.[57] Whoever was posing as Redick was likely a Russian operative.

On August 2, 2017, then-National Security Adviser H.R. McMaster fired Ezra Cohen-Watnick from his position as a top intelligence official on the National Security Council (NSC). Cohen-Watnick was an extremely vocal supporter of Trump, and his dismissal followed the departure of other Trump advocates from the NSC in previous weeks.[58] Later that evening, at least 11 different Twitter accounts posing as Americans—but operated by Russians working for the Internet Research Agency (IRA) in St. Petersburg—tweeted (and retweeted) a message urging that Trump fire McMaster. Among them was the Twitter account @TEN_GOP, which claimed to be the "unofficial Twitter of Tennessee Republicans." This account encouraged its followers to retweet "if you think McMaster needs to go," and many of @TEN_GOP's 140,000 followers were automated "bot" accounts that then automatically retweeted the message. The intended result of this effort was to flood the social media platform with a perception that a groundswell of support was building for the firing of the U.S. National Security Advisor, a former U.S. army general who was (and remains) highly respected by both Republicans and Democrats.

In August 2018, Microsoft disabled six phony websites targeting conservative think tanks and U.S. Senate staff.[59] The sites were apparently designed for a spear-phishing campaign. Another Russian digital influence campaign—dubbed "Operation Secondary Infektion"—used fabricated or altered documents to try to spread false narratives across at least 30 online platforms. According to a report by the Atlantic Council's Digital Forensic Research Lab, and a team of analysts at Facebook who uncovered the operation in June 2019, the network of social media accounts involved "originated in Russia."[60] Similarly, in October 2019 documents from the British government were posted online in an apparent effort to either undermine the ruling Conservative party or sow confusion. According to Ben Nimmo, head of investigations at the social media analytics firm Graphika, this was "either a Russian operation or someone trying hard to look like it."[61]

As detailed in the investigation report by former FBI Director Robert S. Mueller III, online Russian operatives were increasingly active during the 2016 U.S. presidential election,[62] and they have continued to try and influence political issues and debates in America since then. According to a Republican-led Senate Intelligence Committee report released in October 2019, "Russia's targeting of the 2016 U.S. presidential election was part of a broader, sophisticated and ongoing information warfare campaign" using

Facebook, Instagram, YouTube, Twitter, Google, and other major Internet platforms. The report called upon the White House, various government agencies, and the private sector to ramp up efforts to counter this threat in the future. One member of the Committee warned that "Russia is waging an information warfare campaign against the U.S. that didn't start and didn't end with the 2016 election." Another provided a warning: "With the 2020 elections on the horizon, there's no doubt that bad actors will continue to try to weapons the scale and reach of social media platforms to erode public confidence and foster chaos."[63] And in addition to their efforts to influence public perceptions about democratic elections, Russia's information operations also seek to undermine faith and confidence in the independence and legitimacy of other democratic institutions, including courts and the criminal justice system. And by some measures they have succeeded—those who violently attacked the U.S. Capitol Building on January 6, 2021, clearly rejected the legitimacy of over 60 court decisions that upheld electoral results that they did not like.

But beyond Russia, we see other states and non-state actors also using the tools and global connectivity of the Internet—and particularly social media—to launch massive information and disinformation campaigns in order to influence people about politics, science, social norms, and many other (often controversial) topics. For example, in August 2019, Twitter and Facebook revealed a Chinese state-backed information operation launched globally to delegitimize the pro-democracy movement in Hong Kong.[64] Twitter said it had taken down 936 accounts that were "deliberately and specifically attempting to sow political discord in Hong Kong." Facebook said it had found a similar Chinese government-backed operation and deleted fake accounts.[65] Meanwhile, Google shut down 210 channels on YouTube that it said were part of "coordinated influence operations" to post material about the ongoing protests in Hong Kong.[66] Earlier in 2019, Chinese authorities had openly instructed and encouraged hackers to deface websites and attack Telegram accounts of political protestors in Hong Kong.[67] A website hosted on Russian servers, "HK Leaks," posted personal details—names, home addresses, personal telephone numbers—of hundreds of pro-democracy protestors.[68] In their report *Tweeting through the Great Firewall*, the Australian Strategic Policy Institute describes how Chinese language accounts, "leveraging an influence-for-hire network," were used to target Hong Kong citizen and the global Chinese diaspora in a massive effort to discredit the pro-democracy protests.[69]

China has also been using tools of digital influence warfare, like bot and troll accounts, to promote disinformation about—and public debates within—Taiwan.[70] Specific examples have included exposing dissidents' activities, exacerbating political tensions (including a contentious debate over pension payments), and raising suspicions against leading military and political figures.[71] The overall goals of these efforts appear to be to

discredit the secessionist movement, which advocates formal separation from mainland China, and to encourage unity with the People's Republic of China.[72] In addition, China has been aggressively trying to change how the popular information source Wikipedia depicts topics they find politically sensitive.[73] For example, a visitor to Wikipedia would normally see that Taiwan is described as "a state in East Asia." However, anyone can edit Wikipedia entries, and in September 2019, someone (presumably acting on behalf of the Chinese regime) had changed the entry to describe Taiwan as a "province in the People's Republic of China." In the English version of Wikipedia, the Dalai Lama is described as a Tibetan refugee, while the Mandarin version of Wikipedia describes him as a Chinese exile. Similarly, the English entry for the Senkaku Islands said they were "islands in East Asia," but in 2019, the Mandarin equivalent had been changed to add "China's inherent territory." The Chinese Wikipedia describes the 1989 Tiananmen Square protests as "the June 4th incident" to "quell the counterrevolutionary riots."[74] Chapters 2 and 6 of this book will examine China's digital influence efforts in greater detail.

Meanwhile, we also see other authoritarian regimes adopting various strategies and tactics of digital influence warfare. On August 21, 2018, the cybersecurity firm FireEye released a report describing "a suspected influence operation that appears to originate from Iran aimed at audiences in the U.S., U.K., Latin America, and the Middle East. This operation is leveraging a network of inauthentic news sites and clusters of associated accounts across multiple social media platforms to promote political narratives in line with Iranian interests. These narratives include anti-Saudi, anti-Israeli, and pro-Palestinian themes, as well as support for specific U.S. policies favorable to Iran, such as the U.S.-Iran nuclear deal (JCPOA)."[75] Shortly after the report was made public, Facebook announced the removal of 652 users. According to Nathaniel Gleicher, head of Facebook's Cybersecurity Policy, their investigation into a network calling itself "Liberty Front Press" found a direct link to the Iranian government. "For example, one part of the network, 'Quest 4 Truth,' claims to be an independent Iranian media organization, but is in fact linked to Press TV, an English-language news network affiliated with Iranian state media."[76]

A year later, in October 2019, Facebook announced the deletion of 93 Facebook accounts, 17 Pages, and 4 Instagram accounts "for violating our policy against coordinated inauthentic behavior. This activity originated in Iran and focused primarily on the US, and some on French-speaking audiences in North Africa."[77] According to the announcement, "The individuals behind this activity used compromised and fake accounts—some of which had already been disabled by our automated systems—to masquerade as locals, manage their Pages, join Groups and drive people to off-platform domains connected to our previous investigation into the Iran-linked 'Liberty Front Press' and its removal in August 2018."[78]

Facebook also removed 38 Facebook accounts, 6 Pages, 4 Groups, and 10 Instagram accounts that originated in Iran and focused on countries in Latin America, including Venezuela, Brazil, Argentina, Bolivia, Peru, Ecuador, and Mexico. The page administrators and account owners typically represented themselves as locals and used fake accounts to post in groups and manage pages posing as news organizations, as well as directed traffic to other websites.[79] And that same month, Microsoft announced that hackers linked to the Iranian government had targeted an undisclosed U.S. presidential campaign, as well as government officials, media outlets, and prominent expatriate Iranians.[80]

While the above examples are mainly foreign influence efforts—that is, attempts by one country to influence the citizens of other countries—a different sort of case was reported in a recent *New York Times* exposé. Days after Sudanese soldiers massacred pro-democracy demonstrators in Khartoum in June 2019, a digital marketing company in Cairo began deploying what I refer to as "digital influence mercenaries"[81] in a covert operation to praise Sudan's military on social media. The Egyptian company New Waves—run by Amr Hussein, who retired from the Egyptian military in 2001 and describes himself on his Facebook page as a "researcher on Internet wars"—paid new recruits $180 a month to write pro-military messages using fake accounts on Facebook, Twitter, Instagram, and Telegram. On August 1, 2019, Facebook announced that it had shut down hundreds of accounts run by New Waves and an Emirati company with a near-identical name. Facebook said the Egyptian and Emirati companies worked together to manage 361 compromised accounts and pages with a reach of 13.7 million people.[82] According to the report, they spent $167,000 on advertising and used false identities to disguise their role in the operation.[83]

Several other countries are also focused on fabricating narratives that they force-feed to their own citizens through state-owned media and government-controlled bot networks on social media (which we'll examine in chapters 2 and 6). A 2019 report by Oxford University's Computational Propaganda Project found evidence of disinformation and propaganda attempts to manipulate voters and others online in 70 countries.[84] A majority of these attempts were by domestic actors trying to influence domestic targets. For example, Rodrigo Duterte (President of the Philippines) encourages "patriotic trolling" to undermine his critics.[85] A 2017 Oxford Internet Institute report describes how "many of the so-called 'keyboard trolls' hired to spread propaganda for presidential candidate Duterte during the election continue to spread and amplify messages of his policies now that he's in power."[86]

Few states are more committed to spreading disinformation among their own people than Russia. Michiko Kakutani, the cultural critic and author of *The Death of Truth*, has observed how Russia uses propaganda "to distract and exhaust its own people (and increasingly, citizens of foreign

countries), to wear them down through such a profusion of lies that they cease to resist and retreat back into their private lives."[87] A Rand Corporation report called this "the firehose of falsehood"—"an unremitting, high-intensity stream of lies, partial truths, and complete fictions spewed forth with tireless aggression to obfuscate the truth and overwhelm and confuse anyone trying to pay attention."[88] According to the report, "Russian propaganda makes no commitment to objective reality," instead relying on "manufactured sources" and "manufactured evidence (faked photographs, faked on-scene news reporting, staged footage with actors playing victims of manufactured atrocities or crimes). . . . Russian news channels, such as RT and Sputnik News are more like a blend of infotainment and disinformation than fact-checked journalism, though their formats intentionally take the appearance of proper news programs."[89] In fact, as Kakutani notes, "The sheer volume of *dezinformatziya* [see chapters 2 and 6] unleashed by the Russian firehose system . . . tends to overwhelm and numb people while simultaneously defining deviancy down and normalizing the unacceptable. Outrage gives way to outrage fatigue, which gives way to the sort of cynicism and weariness that empowers those disseminating the lies."[90]

Further, monitoring the online activities of its citizens and controlling all forms of access to information online have become hallmarks of authoritarian regimes. For example, Russia has forced search engines to delete certain search results, required messaging services to share encryption keys with security services, and made social network companies store their user data on servers in the country (that presumably they have full access to). Further, beginning in July 2020 Russia will require all smartphones, computers, and smart TV sets sold in the country to come preinstalled with Russian software.[91]

As we'll discuss in chapter 6, sometimes a government will simply shut off the country's Internet access altogether in order to ensure control over what their citizens can say or do online. In November 2019, Iran did this for nearly an entire week. While businesses, universities, government agencies, and other institutions may have suffered from such an act, the underlying logic appeared to be, "Why bother to compete for influence online when you have the power to completely shut off the competition's voices?" Meanwhile, India's shutdown of access to the Internet in Kashmir is the longest ever imposed in a democracy. It began on August 5, 2019, and by mid-December, the province had been without Internet access for 134 days.[92] But thus far, no country has done more than China in using the Internet to influence and control the political and social behavior of its own citizens. With its social credit program and its filtering of Internet search results, China has essentially created its own nationwide digital influence silo (see chapter 5).

Here in the United States, we also see a range of domestic-oriented influence efforts. For example, a number of reports have emerged detailing the microtargeting of veterans in the United States. For years, online scams and fake accounts that exploit or target American veterans have proliferated throughout the Internet, including on Twitter, Facebook, and Instagram.[93] Images of deceased veterans are used as bait in romance scams, memes are spread about desecrated graves in order to provoke anger, and misleading articles about the possible loss of health benefits worry veterans and their families who rely on them. On November 13, 2019, the House Committee on Veterans' Affairs convened a hearing titled "Hijacking Our Heroes: Exploiting Veterans Through Disinformation on Social Media," in which veterans testified about many instances of these things. An extensive report was also published by the Vietnam Veterans of America, describing the "persistent, pervasive, and coordinated online targeting of American service members, veterans, and their families by foreign entities who seek to disrupt American democracy. American veterans and the social-media followers of several congressionally chartered veterans service organizations were specifically targeted."[94]

And on May 22, 2019, an online video of a speech by then-House Speaker Nancy Pelosi was altered to make it seem that her speech was slurred and incoherent and then posted and forwarded by a flurry of Twitter, YouTube, and Facebook accounts.[95] One version of the video, posted by the conservative Facebook page Politics WatchDog, had been viewed more than 2 million times within the first 24 hours of being online and had also been shared more than 45,000 times, garnering over 23,000 comments with users calling her "drunk" and "a babbling mess."[96] The video, as numerous experts in computer science and information technology verified, had been slowed to about 75 percent of its original speed, and the pitch of the speaker's voice had been further modified. A separate video of Pelosi speaking at a news conference was similarly altered (to make her seem like she was stumbling, slurring her words as if she were highly intoxicated)—and then posted to Twitter by then-president of the United States, Donald Trump (see Figure 1.1), amplifying its perceived legitimacy among millions of viewers. In less than 24 hours, the altered video had been viewed more than 3.5 million times on Twitter, earning 70,000 likes and 22,000 retweets.[97]

Of course, modifying the audio and video recordings of politicians in ways that are meant to disparage and embarrass them is nothing new; we've seen that for over half a century. In fact, when a deceptively edited video clip of Democratic presidential candidate Joe Biden circulated on social media in August 2020—which cuts an hour-long speech to less than one minute, retaining only parts of statements and his pauses between words—some observers considered this to be the new norm.[98] Today's

Figure 1.1 Example of then-President Trump using his Twitter account to impugn a political opponent by distributing a doctored video to his followers. (https://twitter.com/realdonaldtrump/status/1131728912835383300)

technology offers cheap and easy ways to do this, along with the ability to distribute the manipulated video as part of a misinformation campaign of unprecedented scale and speed. In May 2019, YouTube took down the altered Pelosi video fairly soon because it has long prohibited altering videos with the purpose of deceiving the public. (YouTube has, however, allowed other hoaxes on the platform so long as they don't promote violence or alter a video clip.) Twitter also kept the video up on its platform. Millions of Trump supporters (many of whom apparently despise anyone who disagrees at all with Trump) tried to influence others' views about Pelosi by distributing copies of this altered video clip and defending its "authenticity" even in light of many news headlines and experts who quickly revealed it to be fake. And unfortunately, as examined in later chapters of this book, there are millions of people in the United States (and billions more worldwide) who believe what they see online regardless of overwhelming evidence proving it's entirely untrue. If nothing else, as Samantha Cole put it, this video "proves that rudimentary editing and willingness to prey on people's hate for public figures is all they need in order to successfully spread misinformation across the Internet."[99]

These are just a few examples of digital influence warfare, provided for purposes of illustration; many more will be provided throughout this book. They reflect how the strategies, tactics, and tools described in chapters 2

and 3 have been used by a variety of states and non-state actors to achieve political power, security, and other kinds of goals and objectives—like social activism, science denial, cyberbullying, economic warfare, and much more. Thus, my approach in this book is to look beyond "political warfare" or information operations and focus instead on the core goal of influencing a target, how it is done, and what enables the influencer to be successful when doing so.

ORGANIZATION OF THE BOOK

Okay, if you've made it this far into the introduction to the book, I'm going to take a leap of faith and assume you want to know more, so here is what the rest of the chapters will cover. First, chapter 2 will focus on the various strategic goals and objectives pursued by these kinds of operations. After briefly reviewing some pre-Internet era examples of influence warfare, we'll discuss a broad range of goals and objectives being pursued through digital influence warfare. For the sake of simplicity, much of this discussion is framed in terms of influencers (or information aggressors) and targets. The means of influencing can vary widely, from spreading blatant lies and disinformation to emotionally provocative (but factually accurate) videos and images, as described in chapter 3. But before choosing which tactics and tools deploy against the target, the influencer must have a clear sense of what goals they want to achieve. We will also examine specific examples like China (with its "Three Warfares" doctrine)[100] and Russia (with its *Information Security Doctrine of 2000*).[101] Both countries have what a Stanford Internet Observatory report calls "full-spectrum propaganda capabilities," and each has amassed prominent Facebook pages and YouTube channels targeting regionalized audiences, though the use of those pages differ according to the kinds of goals and objectives they want to achieve.[102] And this chapter will also review some examples of non-state actors and the strategic goals they pursue using the tactics and tools of digital influence warfare.

Those tactics and tools are the main focus of chapter 3. After a brief explanation of the similarities and differences among the major social media platforms, and the importance of gathering and analyzing quality data on potential targets, the discussion proceeds through three categories of digital influence tactics: deception, provocation, and direct attacks. Within these categories, we find a broad range of specific tactics and terms that may or may not be familiar to most readers. Some tactics are used to discredit institutions that are dedicated to distinguishing between true and false information, while others seek to amplify social grievances, polarization, personal frustration, and anxiety.[103] As we'll see in later chapters of the book, tactics within the categories of deception and provocation are particularly effective for spreading disinformation and for

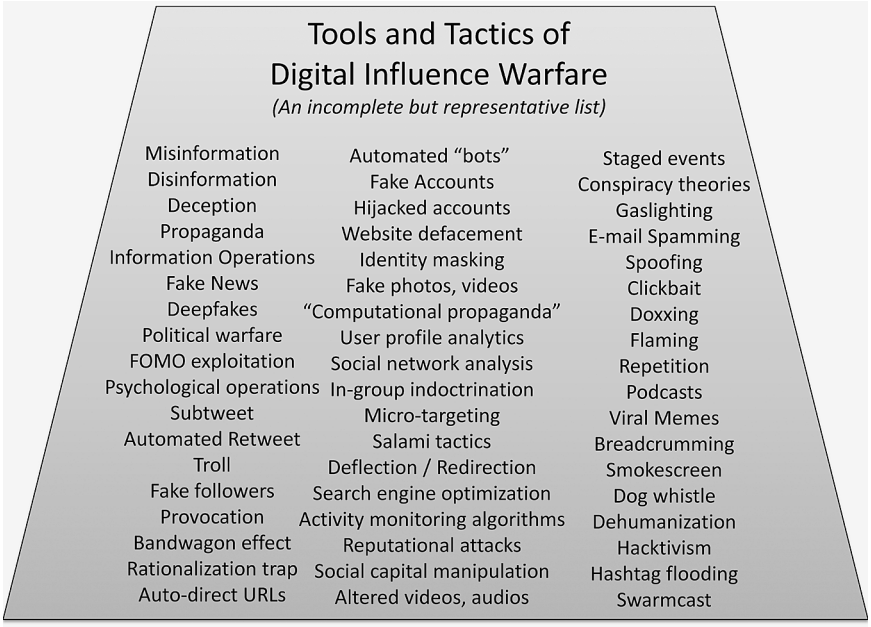

Figure 1.2 An incomplete but representative list of terms used to describe tactics and tools used by digital influence attackers against their targets. For abbreviated definitions of these and many other terms used throughout this book, a "Glossary of Terms" is available online at http://www.DIWbook.com.

exploiting uncertainty and confirmation biases. For the purposes of illustration, Figure 1.2 provides an incomplete but representative list of these tools and tactics.

Some of the terms listed here have been fairly well understood and used for many years to describe various kinds of traditional influence operations—like misinformation (untruths or partial truths mistakenly shared), disinformation (intentional spread of falsehoods), malinformation (leaks and harassment strategies),[104] and propaganda—while other terms (like doxxing, clickbait, hashtag flooding, and search engine optimization) are relatively new and specifically related to the online technologies being used by attackers in various kinds of influence operations. Each of these terms will be explained in this book.

Attempts by influencers to deceive a target are particularly common. As Martin and Shapiro describe in a 2019 report, an influencer can easily produce and disseminate "content that is meant to appear as being produced organically in the target state (e.g., fake accounts, fake press organizations, trolls posting as citizens of the target state)" and may "involve promoting true content as well as false or misleading information."[105] The discussion

in this chapter will travel far beyond the contemporary debates about fake news websites, hijacked social media accounts, and trolls. For example, massive amounts of highly plausible fabricated video and audio material are being disseminated in order to reduce the target society's confidence in a shared reality and intensify their loss of faith in institutions (and each other). In addition to popular social media platforms like Facebook and Twitter, there are many online community sites where users have been prolific in the creation and sharing of deepfake images and videos, "memes," or other viral content. These include Reddit, 4chan, MetaFilter, and Tumblr. Further, recent years have seen the rapid transmission of deepfake technologies from the research lab to user-friendly applications. What started in 2017 with the FakeApp software has now evolved to an open-source version called Faceswap, and the code for accomplishing face swaps is now openly available as a software package called DeepFaceLab.[106] Anyone with even a modest degree of computer literacy can now engage in forms of digital influence warfare.

And both chapters 2 and 3 discuss the importance of assessing the impacts of digital influence efforts and then using this information to refine an influence campaign in order to produce greater levels of effectiveness. This assessment and refinement process reflects the kind of organizational learning we have witnessed thus far among the more sophisticated influence warfare attackers. And the ability to gather and analyze data on the target's reception and reaction to the influence efforts has never been easier, thanks to the Internet. In fact, perhaps the most important contributor to the success of digital influence warfare is the extensive amount of information we provide about ourselves. The advances of social media platforms and search engines have provided a tremendous windfall for the influencers seeking information about you. By monitoring your use of their platforms, companies like Facebook, Twitter, and Google are able to use algorithms to carefully tailor (or "personalize") your online experience in a way they believe you will like, increasing your commitment to using their platforms more frequently.

With a few exceptions, everything you do online can be tracked and recorded, information that can then be used to create a digital profile about your preferences, shopping habits, social and political views, organizational affiliations, religious beliefs, friends, families, and much more. For the influencer, this amount of information about the target is a goldmine, allowing them to fine-tune their strategies and tactics in ways that will be increasingly effective. The websites you visit and the social media posts that you "like" are revealing, but there is much more. The ability to identify whom you "follow" on Twitter or Facebook allows the influencer to gain insight into whom you trust, the people you want to hear from or keep in touch with. Tracking and analyzing the preferences, habits, and so forth of those individuals provides an even greater amount of information

about you that can then further increase the pinpoint accuracy of an influence effort. And from this information, the influencer can then refine and recalibrate their tactics and tools as needed, to include trying different message formats and contents or choosing new targets. Monitoring the success of an influence campaign is made especially easy by social media platforms that capture data on their user's preferences and activities, as described later in chapter 3.

Following these explorations into the strategies and technical mechanics of digital influence warfare, the book turns to examine the psychology of persuasion and influencing. Drawing from research literature on the psychology of persuasion, chapter 4 identifies the main components of influence attempts including the attributes of the influencer, the attributes of the target to be influenced, the content and format of the messages used in the influence attempt, and the context in which those messages will be considered salient and relevant by the target (including how the media often provide that contextual relevance). This discussion is informed by the work of several scholars in the field, including Robert Cialdini's research on the principles of social influence and key concepts like reciprocity, social proof, commitment and consistency, liking, authority, and scarcity.

Chapter 4 also addresses the importance of contextual relevance in effective strategies of persuasion. Because each of us has our own unique collection of beliefs, values, and perceptions about the world and our place within it, the influencer must gather intelligence about their target—including what they like and dislike, what they want to see or don't want to see in their world—and then incorporate that into the strategy and tactics accordingly. You (the target) may not initially care about Topic X, but if the influencer can find ways to directly connect a certain view about Topic X with something you do care about, the odds of you paying attention will naturally increase. So, this chapter explores how contextual relevance can be identified and manipulated by an information aggressor in order to increase the target's likelihood of being influenced. Information that is conveyed via a trusted source, or that is novel, sensational, or emotional (for example), can attract the attention of the target more than other kinds of information. Here's an example of why this matters: in order to effectively utilize the tactic of "ragebait" to provoke outrage among the target audience, it is first necessary to understand the values and beliefs that frame what that audience views as something to get worked up about.

Finally, the chapter briefly reviews how uncertainty, fear, conformity, and confirmation bias play significant roles in how people respond to these influence efforts. There is a considerable body of research in social psychology that identifies how group identity impacts an individual's information processing choices. For example, consider the following two statements: "I hate X, and something bad is being reported

about X; therefore, it must be true." And "I like A, and something bad is being reported about A; therefore, it must be lies." This is essentially how in-group identity and "othering" can shape perceptions of a target audience. As a recent Rand Corporation report notes, influence operations "are likely to take the form of targeting subsections of the population to intensify divisions and polarization rather than attempting to shift or create new beliefs wholesale in a population."[107] These subsections are manifest in the ways that a majority of Americans now access their news primarily via social media.[108]

In many cases, the target's values and beliefs are shared among groups of people, allowing for various forms of collective behavior in areas of politics, religion, social norms, and so forth. This is particularly the case inside what I call digital influence silos, the topic of chapter 5. Essentially, these are virtual bubbles of information in which we surround ourselves with factoids and narratives that confirm what we want to believe and effectively block out any kinds of information that questions or contradicts those beliefs. Many readers will have seen these referred to by other terms as well—like filter bubbles, influence bubbles, and echo chambers[109]—but the overall concept is that we live in a world in which we are able to surround ourselves extensively and exclusively with information that confirms what we want to believe. Our ability to avoid (or even ignore the existence of) information that questions or contradicts what we want to believe has never been greater. This gives a huge advantage to the malicious actor using digital influence strategies and tactics to achieve their goals. All that is now required is to tailor their message in a way that conforms to what we want to believe.

For example, if the goal is to sow discord and animosity among members of a community, the data and tools are now available for identifying disagreements and seams of latent distrust. Within any community, there are always in-groups and out-groups, and it is increasingly easy to identify the members of each. The next step is to simply provide the kinds of information (or disinformation) within each influence silo that will exacerbate distrust toward those "others" outside the silo, add more fuel of hostility to the disagreements, and then sit back and watch them go at each other's throats. If the goal is to encourage a "rally around me and my cause" effect, first convince the members of a particular influence silo that your issue is directly tied to whatever they care most about. Then create the illusion that what they care most about is threatened by those outside their influence silo and that what you want is also threatened by the same out-group. You can also deploy tactics of identity politics to convince some members of the influence silo that they are not fully committed or being true to the in-group unless they act in some way (usually in defense of what you want them to believe is threatened). As Oxford University's Philip Howard notes, digital influence efforts are most effective when the

messages are "delivered by a relatively enclosed network of other accounts and other content that affirms and reinforces what people are seeing."[110]

As we'll address in chapter 5, modern social media platforms have helped create a variety of competing influence silos in at least two ways.[111] First, we have a tremendous amount of freedom to actively seek out information that we believe will provide us with what we want to know, and quite often, we prefer sources of information that confirm what we already tend to believe. Further, we can ignore (or even block out) other information sources that question or contradict what we want to believe. Meanwhile, influence silos are also nurtured by an automated personalization of information involving algorithms that we have virtually no control over and yet show us what the computer believes we will want to see. As a result, the computer becomes a channel through which we can be effectively influenced, and an increasing number of malicious actors are hacking into the channel to manipulate what we see and hear. This is why the term "digital" is used throughout the book to mean something that has direct relation to the online information ecosystem.

In chapter 6, we examine the ultimate manifestation of digital influence silos, in two separate forms. In authoritarian countries, as described in the first part of the chapter, we find governments seeking (and sometimes successfully achieving) a form of information dominance in which all the information available to the country's citizens is highly controlled. In one sense, they are able to establish a nationwide influence silo, within which individuals only see and hear information that has been preselected for their consumption. When this is possible, your target audience has no choice but to hear your narrative, and yours alone. Authoritarian regimes are perennially manipulating public opinion and perceptions by curating and controlling the information their population is allowed to see, and the Internet provides the means to do this in more ways than we have ever known before. Further, by increasing uncertainty among members of the target audience, within an environment where your ability to influence the target is unchallenged and unconstrained, you can convince them of virtually anything.

Authoritarian regimes also seek to establish and maintain information dominance as a means of reducing uncertainty among its citizens about what to believe. Essentially, they replace one choice with another: trust what we tell you and do what we tell you or face dire consequences. Countries like China, Russia, Turkey, North Korea, and Iran respond to questions of uncertainty by simply imposing their own information dominance. Disinformation is crafted by government leaders and fed to the population through government-controlled media and multiple online forms of communication. Competing sources of information, particularly if they question or contradict official narratives, are simply banned; uncontrollable journalists are jailed (and in some instances killed); search

engine results are limited to only acceptable sources of information; and many other tactics are used to create informational barriers and to shape perceptions of the population.

However, this kind of information dominance is not readily available in truly democratic countries, where freedom of speech and expression is protected and where citizens can access a broad range of information sources. So instead, influencers wishing to achieve a similar level of power will often pursue what I call "attention dominance," something that is made increasingly possible through the algorithms of social media platforms, search engines, and website trackers. We'll discuss examples of this in the second part of chapter 6. Replicating information dominance may be more difficult in an open democratic society, but the Internet has now provided the means to do so in unprecedented ways. By crafting and co-opting influence silos that reinforce the cognitive biases and preferences described in the previous chapter, the tools and tactics of digital influence warfare can be used to shape reality in a way that makes it increasingly easy to block out any dissenting types of information and their sources.

Platforms like Facebook and Twitter use a variety of algorithms to filter the information we see in our daily "news" feed based on what we most likely "want" to read. The business model of social media companies relies on clicks and preferences, not telling people what they "should" know. This, in turn, leads to a fragmented digital information environment that can virtually isolate people within ideologically partisan communities that have no access to (or interest in) any type of information that does not conform with their preferences, prejudices, and beliefs. Because of this kind of environment, the behaviors of the target can be manipulated, particularly by influencers who appear to be aligned with the target's previously established preferences. As a result, we are now rapidly hurtling toward a future of influence attacks using increasingly realistic deepfake videos, audio clips, and images—many of which are intentionally trying to ruin trust in specific individuals and institutions.[112] Together, these factors explain how blatant lies and disinformation can incentivize behavior like clinging to beliefs even in the face of overwhelming evidence that proves those beliefs to be in error. Given the choice between an inconvenient truth and an enjoyably confirming falsehood, many of us will choose the latter.

Finally, the book concludes with a brief look at what we should anticipate for the future of digital influence warfare. For example, advances in artificial intelligence will make fake photos and videos increasingly difficult to detect. Russia will increasingly rely on its complex international network of hackers, activists, and informal propagandists to further pursue its strategic and foreign policy objectives, while China will expand its use of use of Chinese citizens and ethnic Chinese abroad to further its control over key narratives. These and other countries will find new ways to mask their involvement in digital influence efforts against domestic and foreign

targets.[113] Meanwhile—as explained by a recently published Rand Corporation report—"the conflicts for ideological supremacy emerging between influence silos are encouraging new forms of widespread cyber-harassment, and in time this will result in the Internet becoming a notably crueler and more intimidating space."[114] The future will likely bring an increase in various forms of cyber harassment attacks, such as creating false websites with allegedly compromising information; generating fake videos using high-grade digital mimicry programs that allegedly show the targets stealing, killing, or in intimate contexts; hacking official databases to corrupt the targets' tax or police records; sending critical, crude, and self-incriminating emails to dozens of friends and colleagues, seemingly from the target, using spoofing techniques to conceal the origins of the messages; and hacking targets' social media accounts in order to post offensive material supposedly in their name.[115]

CONCLUSION

As Carl Miller explains, "With states, political parties and individuals jockeying for ever-greater influence online, you and your clicks are now the front line in the information war."[116] Whether you are using your smartphone, desktop, laptop, tablet, Internet-connected television, or any other means to go online, anything you see on that screen is inherently digital—the words, images, and sounds are all based on various compositions of digits (1s and 0s). This technological environment (and the ways in which we interact with it) offers new tools and tactics for influence warfare. The revolution in communications technology driven by the Internet has created a new, more expansive market of ideas. Individuals are now empowered to reach massive audiences with unfiltered messages in increasingly compelling and provocative packaging, rendering the competition for mass influence more complex. The emergence of new means of communication and new styles of virtual social interaction have transformed the context for mass persuasion and have expanded opportunities for anyone to disseminate their message.[117] Social media is particularly central to digital influence warfare. Not only can we easily become overwhelmed by the volume and diversity of information available in our social media account feeds but also much more of that information is trivial, one-sided, and fake than we've ever encountered before. This makes it increasingly difficult to distinguish fact from fiction, or evidence-based arguments from biased opinion, and the result is greater uncertainty and misperceptions about what is true and what is not.

This book is not just about technology—automated fake accounts, data algorithms, deepfake videos, and so forth—that underpins digital influence warfare. The book is essentially about real people doing real things with

real consequences. It's about the intersection of human behavior, beliefs, technology, and power. Whether they're trolls paid by Russian government agencies, politicians who lie, Chinese censors, or violent extremists, the Internet is simply the means by which they are trying to achieve certain influence objectives. Further, many forms of digital influence increasingly focus on getting real people—our family, friends, colleagues—to share and retweet lies and disinformation, something that has become all too easy today given the rise of digital influence silos and our own reliance on cognitive biases to sort through a confusing avalanche of information. Unfortunately, democratic societies are considerably vulnerable to disinformation, resulting in distorted public perceptions fueled by algorithms that were originally built for viral advertising and user engagement on Internet platforms. Further, as Thomas Rid observes, "Disinformation corrodes the foundations of liberal democracy, our ability to assess facts on their merits and to self-correct accordingly."[118] Today's disinformation can include a wide variety of digital items—from images and videos to official documents—that can all be fabricated or altered in ways that manipulate our perceptions and beliefs about something. And disinformation is just one of several variants of digital influence warfare. Leaking confidential documents and correspondence to the public for malicious effect isn't considered disinformation (the spread of falsehoods), but it is an act driven by similar kinds of influence strategies and goals. Similarly, factually accurate information can be used to provoke certain kinds of behaviors among the target audience of an influence operation.

To sum up, a wide array of strategies, tactics, and tools of digital influence warfare will increasingly be used by foreign and domestic actors to manipulate our perceptions in ways that will negatively affect us. According to a UNESCO report, the danger we face in the future is "the development of an 'arms race' of national and international disinformation spread through partisan 'news' organizations and social media channels, polluting the information environment for all sides"[119] Tomorrow's disinformation and perceptions manipulation will be much worse than what we are dealing with now. The future also promises to bring darker silos of deeper animosity toward specifically defined "others" who will be deemed at fault for the grievances of the silo's members. With this will come a higher likelihood of violence, fueled by emotionally provoking fake images and disinformation (from internal and external sources) targeting the beliefs of the silo's members. This is the future that the enemies of America's peace and prosperity want to engineer. We must find ways to prevent them from succeeding. At the end of the day, one of my goals for writing this book has been to encourage each of us to look more closely at how our own decision-making is being influenced each day, by whom, and what their goals might be. When we stop and think about the influences

we are experiencing, we tend not to be as open to digital influence and exploitation. If we don't treat this battlefield with greater levels of attention and urgency, identifying and confronting the various forms of digital influence warfare used on that battlefield, we will succumb to whatever our adversaries' strategic goals are. We must confront this, collectively and urgently.