Preface

(Submitted Draft, September 4, 2020)

Let me begin with an honest self-reflection. I have published over 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; a highly polarizing presidential election; and much more). 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 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 a particular focus on how websites, blogs, e-mail, 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 U.K. and the U.S. But an increasing number of other countries are also engaged in similar activity, 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 live-streaming video 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 states 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 understood 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 literature 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 targets" believe or think they know is based on false information, leading to a sense of superiority over those naïve "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), emotional provocation, and outright attacking the target (to include 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 (outrange, 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 new studies about what I now call digital influence mercenaries, and found many examples of non-state actors who are profiting by deceiving and provoking people on social media. My journey also led me to the research on technological tools used by states 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). However, there is so far limited agreement

on what should be done, or on who should do what, and even less agreement on how each of us can play an important role in this response. So, in the concluding chapter I briefly explore what governments, the private sector, and individual citizens can do to confront digital influence warfare efforts today and into the future. The first conclusion I arrived at is that we—the targets—must recognize what is going on. When we stop and think about the influencers behind the information we see and hear, we tend not to be as open to exploitation. Second, 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 can put pressure on private sector media firms to combat and counter-attack those seeking to spread disinformation. Facebook's current policy of allowing political advertisements to include bold-faced lies is unhelpful in this regard. 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 cyber attack that requires more than network systems firewall and security. Confronting and deflecting these digital influence efforts requires 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. Absent that, our enemies will continue trying.

While this book was being written, dozens of election campaigns were launched by a flurry of Democratic candidates hoping to be the one to run against (and defeat) the incumbent in November 2020. Now that field of candidates has been narrowed to just one Democratic nominee, and we have already seen many examples of how the strategies and tools of digital influence warfare can be used against political opponents. We also endured various forms of social mediated disinformation, disorientation and conspiracy theories as the deadly COVID-19 virus spread to countries around the world. In 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, some of the analysis 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. To be honest, I have written and re-written 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 co-workers or guest lectured in my courses, and even co-authored 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 work of these people most enlightening, particularly the hard-working folks at the Oxford Internet Institute's Computational Propaganda Project, 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.

I also want to include here a special shout-out to professors and mentors in graduate school many years ago who took me under their wing and modeled for how to 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 9 years there, particularly from my friends and colleagues in the Department of Social Sciences and the Combating Terrorism Center, as well as from 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.