Chapter 1

Retweets to Midnight:

Assessing the Effects of the Information Ecosystem on Crisis Decision Making between Nuclear Weapons States

Danielle Jablanski, Herbert S. Lin, and Harold A. Trinkunas

What if the Cuban Missile Crisis had taken place in today's global information environment, characterized by the emergence of social media as a major force amplifying the effects of information on both leaders and citizens? President Kennedy might not have had days to deliberate with the Executive Committee of the National Security Council before delivering a measured speech announcing to the world the discovery of Soviet medium- and intermediate-range nuclear-armed missiles in Cuba.[1]

Nongovernmental open source intelligence organizations like Bellingcat could have used commercially available satellite imagery to detect the presence of these missiles and publicize them to the world on October 12, 1962, four days earlier than the president did. Imagine pictures of the missile sites going viral on social media, alarming millions around the world. Imagine that these real-time images were accompanied by deliberate information operations from adversaries seeking to cast doubt on the facts to sow confusion and cause paralysis among domestic populations and between NATO leaders, as well as by internet trolls promoting misinformation and reposting and propagating tailored information leaks.

The shooting down of a U-2 spy plane over Cuba might have been news within the hour, becoming the subject of numerous tweets and relentless commentary on Facebook and other platforms. When the Joint Chiefs of Staff's recommendation to invade Cuba was overruled

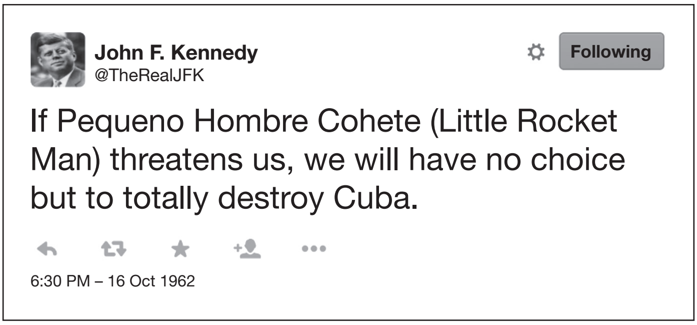

John F. Kennedy
@TheRealJFK

If Pequeno Hombre Cohete (Little Rocket Man) threatens us, we will have no choice but to totally destroy Cuba.

6:30 PM – 16 Oct 1962

Figure 1.1. Hypothetical tweet by President John F. Kennedy during Cuban Missile Crisis.

Source: Scott Sagan, “The Cuban Missile Crisis in the Age of Twitter,” lecture at Stanford’s Center for International Security and Cooperation, April 3, 2018.

by President Kennedy, “alt” social media accounts that served as fronts for disgruntled Pentagon officials might have leaked the proposed invasion plan to induce the administration to reverse course on the chosen alternative—a blockade. Pressured by public opinion and the Joint Chiefs of Staff, President Kennedy might not have had the luxury of picking which of Premier Khrushchev’s letters to respond to, which the historical record shows helped to de-escalate the crisis. In this situation, which former secretary of defense William J. Perry has characterized as the closest the world has come to nuclear catastrophe, the current global information ecosystem could have magnified the risk of the conflict’s escalating into all-out nuclear war.[2]

What’s New? Characteristics of the Modern Information Ecosystem

Social media and the resulting dynamics for interpersonal interconnectivity have increased the volume and velocity of communication by orders of magnitude in the past decade. More information reaches more people in more places than ever before. Algorithms and business models based on advertising principles utilize troves of user data to draw aggregate inferences, which allow for microsegmentation of audiences and direct targeting of disinformation or misinformation. Mainstream-media outlets no longer serve their traditional role as gatekeepers with near-universal credibility.[3] In this ecosystem, propaganda can rapidly spread far and wide, while efforts to correct false information are more expensive, often fall short, and frequently fail altogether.

Nor are all of the voices on social media authentic. Some inauthentic voices are those of paid human trolls, for example from the Internet Research Agency, revealed to have created and spread false information on behalf of the Russian government prior to the 2016 US presidential election.[4] Others are Macedonian entrepreneurs who at one point discovered ways to monetize an affinity among some voters for fake news critical of Hillary Clinton.[5] Some voices are not even human, as demonstrated by the introduction of "bots"—automated social media accounts designed to mimic human behavior online that further complicate our ability to discern fact from fiction within the ecosystem.

Rapid transmission of content and curated affinity networks polarize citizens around divisive issues and create waves of public opinion that can pressure leaders.[6] So many different narratives emerge around complex events that polities splinter into their disparate informational universes, unable to agree on an underlying reality. Does this unprecedented availability of information and connectivity amplify the ability of actors to sow discord in the minds of the domestic publics and even the leadership of adversaries? Could these dynamics affect leaders and citizens to the degree that miscalculation or misperception can produce crisis instability ultimately leading to a nuclear exchange? Can governance mechanisms be designed and implemented that are capable of countering and combating the manipulation of information in this ecosystem?

This volume argues that the present information ecosystem increasingly poses risks for crisis stability. Manipulated information, either artificially constructed or adopted by a strong grassroots base, can be used by interested actors to generate pressure from various constituencies on leaders to act. At the same time, these leaders themselves face information overload and their ability to distinguish between true and false information may be impaired, especially if they are receiving information simultaneously from their own sources and other sources from within their constituencies. Such confusion can ultimately lead to inaction or bad decisions. Or, this environment might produce an accelerated reaction based on slanted or unanalyzed information. Most worrisome is the possibility that the rapid spread of disinformation or misinformation via social media may in the end distort the decision-making calculus of leaders during a crisis and thereby contribute to crisis instability in future conflicts, the effects of which could be most severe for nuclear weapons states.

The Psychology of Complex Decision Making and Nuclear Crisis

Many theories of deterrence rely on the rationality assumption, namely that a rational actor can be convinced that the cost-benefit ratio associated with initiating an attack is unfavorable due to a credible threat of retaliation by the adversary. The risk of a nuclear exchange during the Cold War led theorists to focus on how leaders might approach crises and what could be done to avert deterrence failure. This prompted debates about a range of putatively rational actions that nuclear states might engage in to build a reliable framework for deterrence: reassurances to allies by extending the nuclear umbrella, force postures designed to ensure a survivable retaliatory capability, credible signaling to convince adversaries that any attack would meet with massive retaliation, etc.[7]

But human decision makers are just that—human—and a great deal of psychological research in the past few decades has demonstrated the limits of rational thinking and decision making. Paul Slovic has written extensively about the human brain, decision making, and limits for comprehending the weight of decisions that could imperil large numbers of human lives. Various psychological processes come into play when considering a cognitive calculation on the value of lives lost in large numbers, including psychic numbing, tribalism, the prominence effect, imperative thinking, and victim blaming. As Slovic and Herbert Lin argue in chapter 3, this implies that leaders facing the task of making a decision on whether to order the use of nuclear weapons find it difficult to operate "rationally."

Psychology also tells us that—more often than not—fast, intuitive judgements take precedence over slower, more analytical thinking. Fast thinking (also identified as System 1 thinking by the originator of the concept, Daniel Kahneman) is intuitive and heuristic, generating rapid, reflexive responses to various situations and—more often than not—useful in daily life. Slow thinking (also known by cognitive psychologists as System 2 thinking) is more conceptual and deliberative.[8] Although both are useful in their appropriate roles, their operation in today's information ecosystem can be problematic. "Fast thinking is problematic when when we are trying to understand how to respond to large-scale human crises, with catastrophic consequences," Slovic and Lin write. "Slow thinking, too, can be incoherent in the sense that subtle influences—such as unstated, unconscious, or implicitly held attitudes—can lead to considered decisions that violate one's strongly held values." The prevalence of heuristic and "imperative thinking" among humans suggests that an overarching important goal, such as national defense in the face of a nuclear crisis, would likely eclipse consideration of second-order effects and consequences, such as the likelihood of massive loss of life on all sides or catastrophic effects on the global environment, to the extent that such discussion is actively, if not subconsciously, avoided.[9]

Observers have always anticipated that leaders would be under severe time pressures when deciding whether or not to use nuclear weapons, the most important of which is "launch on warning," the pressure to launch fixed land-based ICBMs before they can be destroyed on the ground by incoming enemy warheads. Fast, reflexive thinking (i.e., System 1 thinking) is more likely to be used under the kind of pressure this scenario highlights. Against a ticking clock, combined with the difficulty of comprehending the consequences of nuclear conflict, the argument that rational and deliberate decision making and deterrence will likely prevail, particularly under the added weight of the misinformation and disinformation that might propagate through the global information ecosystem during a crisis, is a highly debatable proposition.

The possibility that decision makers may rely on incorrect perceptions of potential adversaries has long been an important critique of rational deterrence theory. International relations theorists such as Robert Jervis have argued that the failure of deterrence can frequently be attributed to misperception among leaders: of intentions, of capabilities, of the consequences of conflict, etc. This misperception can have its roots in leaders' psychology, in lack of information, and in leaders' assumptions about what information the other side has or how they in turn perceive the situation.[10]

In the 1980s, Jervis had already argued that misperception was a quite common cause for deterrence failure. In today's global information ecosystem, there are more data available than ever before. But rather than reducing the likelihood of misperception through the greater availability of information about potential adversaries, the present information environment provides unprecedented opportunities for manipulation of leaders' and publics' perceptions about intentions, capabilities, and consequences of conflicts—cheaply, rapidly, and at scale.

Tools and Tactics in the Modern Information Ecosystem

Social media have emerged as a modern vehicle for changing narratives. Social media are arguably optimized to try to keep users in a "fast" pattern of thinking, promoting impulsive and intuitive responses to engage users emotionally and maximize both advertising revenue and user experience.[11] This characteristic of social media platforms may also provide avenues by which these same users can be manipulated more effectively for political aims. Although the ability for propaganda to be both insidious and anonymous is not a new phenomenon, automation, algorithms, and big data are being employed by various actors to selectively amplify or suppress information viewed by hundreds of millions of people via social media and online networks.[12] There is evidence of targeted influence campaigns in at least forty-eight countries to date.[13] Facebook, YouTube, WhatsApp, Instagram, and Reddit have also been platforms for a variety of divisive information operations.

As Mark Kumleben and Samuel C. Woolley note in chapter 4, campaigns have often made use of networks made of "bots"—partially or wholly automated—to introduce and circulate false or malign information, craft and control the narrative at the outset of a real event, and depict a manufactured consensus base around an issue. For example, an estimated 15 percent of Twitter's approximately 335 million users (as of 2018) are bots. Bots are employed as a tool to promote a mix of authentic and inauthentic content, automate the amplification of specific sources, disrupt and overwhelm channels and conversations with irrelevant noise, and harass individuals or groups online to silence or intimidate them.

Information operations have more than commercial or political/electoral implications. We are also witnessing an increase in states using such strategies to shape potential battlefields. Using the example of information operations against NATO, Kate Starbird employs a mixed-method analysis in chapter 5 of this volume to understand online communities and their communication patterns. In this case study, Starbird used a data set that included 1,353,620 tweets (75 percent retweets) from 513,285 sources. She mapped accounts into five clusters based on narrative and user characteristics and looked at the interactions between them. One cluster, associated with NATO and other verified accounts, carried positive information about anniversaries and anecdotes that promoted support for NATO. Another cluster, which she named the "international left, anti-NATO," involved content related to past negative NATO actions and events. Pro-Russian and official Russian accounts deliberately penetrated, amplified, and mingled with this cluster and a third cluster that she characterized as pro-Trump and alt-right. Starbird's team also observed interactions between the mainstream-media cluster and the official NATO clusters, both highly critical of Russia. Her findings illustrate structured narratives being shaped, infiltrated, and leveraged by different actors and audiences online in real time.

Starbird's chapter is a case study in how networks and communities can converge and diverge on the same topic, with various actors rescoping and redirecting traffic and conversation to achieve political effects. These dynamics may be an important factor in shaping support for NATO among the general public, which may ultimately have important security implications. A broader lesson is that analysis of information operations often focuses on election meddling while ignoring potentially broader campaigns with the longer-term goal of undermining partnerships, alliances, and international stability. Such operations may, for example, be used by adversaries to propagate social media statements by American leaders in combination with "fake news" and disinformation in an effort to affect the perceptions of traditional US allies—NATO member countries, South Korea, or Japan—of the reliability of US treaty commitments or extended nuclear deterrence.

Most people would like to believe that they are hard to fool, but it is often difficult to distinguish between authentic and fake content. Research shows that deliberately false information is often deeply embedded in otherwise factual content. Text and image repetition plays to the "illusory truth" bias (i.e., information seen more frequently is perceived as being more likely to be true). Discernment of bad information can depend on individual levels of education and skepticism. Repetition of a concept, even by attempting to dispel a false narrative, can serve to reinforce original beliefs. Compatibility with worldview, coherence of argument, credibility of source, consensus of others, and supporting evidence offer a scaffolding framework that supports what individuals come to believe. With decreased trust in media, conflicting expert opinions around every topic, and selective and simplistic sorting of complex information, social media platforms provide a growing medium for manipulation of individuals, groups, and larger institutions.

The current information ecosystem may very well be experiencing a "post-truth" moment. Despite traditional Western enlightenment thinking that it was only possible to discern truth based on evidence and science, mass susceptibility to media environments where facts are less appealing than emotion and opinion dates back centuries. Rose McDermott writes in chapter 2, “Many people rely on their emotions as the most readily accessible, accurate, and immediate source of truth.” They prefer fast, emotional processing because “analysis of abstract knowledge requires so much additional effort.” Appealing to this cognitive trait eases information overload and allows us to fall back on biases—biases that “make us prone to systematic error or susceptible to systematic manipulation by others.” McDermott also highlights the tendency for individuals to more readily accept information that aligns with previously held convictions.

Extrapolating from the psychological findings described above to a systematic impact on public opinion from social media is suggestive, but not dispositive. Rather than social media altering public opinion, it may be that social media reveal more clearly the full spectrum of opinion. Exposure to propaganda may not result in conversion or persuasion to new and different points of view, though it may crystallize or harden existing prejudices and political inclinations. Indeed, it is a fact that polarization and tribalism predate social media, as do interdependent narratives formed about systems, identities, and issues. Ben O’Loughlin argues in chapter 9 that certain narratives or worldviews have a stronger role in shaping issue-specific narratives, such as views on nuclear weapons and war. He also argues that “through identity narratives, communication is used to shape behavior.” In his view, it would be difficult for disinformation and misinformation to penetrate long-standing social networks, although it could be a very powerful tool for mobilization.

Crisis Stability and Escalation Risks

Although the underlying causal mechanisms need additional study, the present information environment—where misinformation and disinformation can be used to potentially alter, polarize, or harden leaders' and publics' perceptions of intentions, capabilities, risks, and consequences during international disputes—raises the specter that stability during crises can be deliberately manipulated, at greater speed, on a larger scale, and at a lower cost, than at any previous time in history.

In chapter 7, Kristin Ven Bruusgaard and Jaclyn Kerr argue that "if there were ever to be perfect crisis stability—and thus a perfect absence of risk of a crisis resulting in nuclear use—then nuclear weapons would serve no useful purpose as a deterrent and, in the absence of other forms of deterrence, aggressive states would have no reason to fear undertaking aggressive subnuclear military action." But there is no perfect crisis stability, and today's global information ecosystem makes the problem worse rather than better. Crisis stability is influenced by a battle of perceptions in which each actor determines whether it is in its interest to strike first or to take actions which knowingly could be perceived as escalatory. Bad information, public disorder, and panic can all lead to miscalculations or misperceptions that can prompt a state to strike first or escalate.

Kelly Greenhill argues in chapter 6 that the current information ecosystem contributes to political escalation through the "nonmilitary shifts in scope and intensity whereby states or actors adopt more aggressive rhetoric, articulate more expansive war aims, or announce decisions to relax or otherwise shift the prevailing rules of engagement." In this ecosystem, rumors, conspiracy theories, myths, propaganda, and fake news—what might be called "extra-factual information"—can inadvertently catalyze resolve that did not previously exist in leaders, the public, and adversaries, or can tie leaders' hands once escalatory rhetoric is broadcast publicly.

Effective signaling "requires reducing risks of misperception and tailoring messages accordingly," according to King's College London's Heather Williams.[14] But the use of social media by senior political leaders can increase the risk of misperception. For example, official and personal social media have been used by the US president both to diminish and to woo partners and adversaries and to announce presidential policy intentions. Military leaders have used social media to announce new capabilities. Adversaries, whether government or military officials, in both their individual and official capacities, monitor social media messaging of leaders and their subordinates. This means that leaders and the bureaucratic institutions that support them are vulnerable to information operations because they monitor social media to understand, at least partially, leaders' intentions.

Moreover, it is entirely possible that actors peddling falsified information about an adversary can come to believe their own propaganda. In chapter 8, Jeffrey Lewis outlines examples of how various false stories propagated by Russian official sources related to US missile defense systems in Europe affected arms control negotiations: "The [Russian] disinformation campaign was part of a continuing effort to paint US missile defense systems to be deployed in Poland and Romania as systems that could be converted to house offensive missiles, armed with nuclear weapons, and used to decapitate the Russian leadership." While completely untrue, Russian arms control negotiators came to believe this to be a reality, insisting that these false stories were evidence of American breaches of existing arms control agreements.

In addition, deliberate interference using information operations during a time of crisis could accompany a physical attack. Ven Bruusgaard and Kerr identify three overarching effects related to manipulating information and narratives on nuclear weapons: to enhance deterrence by confusing adversaries' understanding of your capabilities or their own capabilities; confusing their understanding of your goals and intentions; and instigating conflict by a third-party proxy with manipulated or false information sparking political, military, or public crises. And as Kumleben and Woolley point out in chapter 4, even before an attack, long-term information operations and manipulation of false alarms could lead to desensitization, false evacuations, and countermanding of real warnings. The very rapid dynamic of information flows on social media affects the two main principles for crisis stability articulated by historian Lawrence Freedman: the need for government leaders to outline clear objectives *before* a crisis and the desire for *increased* time for processing information during a crisis.[15] The dynamics of the global information ecosystem and their real-time effects on crises can no longer be ignored.

Conclusions

Against a backdrop of multiplatform communication suffused with a mix of information—true and false, official and unofficial, from friend and from foe, emotionally charged and serenely rational—the dynamics of the modern information ecosystem suggest unprecedented pressures on government decision makers during crisis. The timelines for decision making will be far more constrained, a fact likely to lead to a greater reliance on fast thinking by decision makers just at a time—during a crisis—when slow, deliberate, analytical thinking is most important.

As Ven Bruusgaard and Kerr write in chapter 7, information plays a critical role in assessing the actions, motives, and likely responses of other states, and thus informs available options for decision-making, reinforces or undermines biases, confuses or clarifies analysis, and dampens or amplifies pressures through public feedback or panic.[16] During the Cuban Missile Crisis, analog systems, government monopoly on both secret and public information, and the relatively lengthy decision time available to leaders helped to avoid disaster. Kennedy and Khrushchev had many days to deliberate and shape their proposed responses before they were revealed to the public.[17]

Leaders temperamentally reliant on fast thinking may well neglect or undervalue important considerations that should be taken into account. Indeed, if today's media environment existed during the Cuban Missile Crisis, American and possibly Soviet leaders would have been awash in a sea of unverified or unverifiable, emotionally laden, and politically fraught information on their adversary's intentions, force postures, and public opinion. Decision makers predisposed to impulsive thinking and whose personal information environments will include raw social media feeds as well as vetted information from their support agencies may well be more likely to act on those impulses.

The emergence of a crisis will also affect the modern information environment for the average person. Data show that publics strongly disapprove of inaction in response to provocation.[18] The simplistic, expedient, and repetitive nature of modern social media communication is likely to reinforce this sentiment, driving more citizens into a greater reliance on fast, reflexive thinking of their own. The result could well be increased public pressure on leaders to act more quickly than would be wise.

As for the future, it is reasonable to expect that social media platforms will continue to grow and evolve in this "new normal" ecosystem and provide ever greater ease for like-minded individuals to connect and share information, both innocent and malign. New tools—including video and audio deep fakes, explicit coordination between information operations and real-world events, and microtargeting of individuals with customized messaging—will enable further pollution of the information ecosystem. Dissecting the effects and impacts of information operations on publics and adversaries prior to a crisis, on publics and leaders during a crisis, and ultimately on decision making will be increasingly difficult in this rapidly evolving environment. Introducing, framing, and controlling narratives has become a new type of warfare being fought online each day—often in 280 characters or fewer.

Notes

[1] Len Scott and Steve Smith, "Lessons of October: Historians, Political Scientists, Policy-Makers and the Cuban Missile Crisis," *International Affairs* 70, no. 4 (October 1994): 659–84, https://doi.org/10.2307/2624552.

[2] Former secretary Perry has said the odds of nuclear war were worse than those estimated by President Kennedy, which was "somewhere between one out of three and even," because Kennedy did not know at the time that the Soviets already had tactical nuclear weapons on Cuba and authorization to use them in the event of a US invasion (which was the unanimous recommendation of the

Joint Chiefs of Staff at the time). William J. Perry, *My Journey at the Nuclear Brink* (Stanford, CA: Stanford University Press, 2015).

3. Philip M. Napoli, "Social Media and the Public Interest: Governance of News Platforms in the Realm of Individual and Algorithmic Gatekeepers," *Telecommunications Policy* 39, no. 9 (October 2015): 751–60, https://doi.org/10.1016/j.telpol.2014.12.003; Michael Latzer, Katharina Hollnbuchner, Natascha Just, and Florian Saurwein, "The Economics of Algorithmic Selection on the Internet," in *Handbook on the Economics of the Internet*, ed. Johannes M. Bauer and Michael Latzer (Cheltenham, UK: Edward Elgar, 2016), 395.

4. Neil MacFarquhar, "Inside the Russian Troll Factory: Zombies and a Breakneck Pace," *New York Times*, February 18, 2018, https://www.nytimes.com/2018/02/18/world/europe/russia-troll-factory.html.

5. Samanth Subramanian, "Inside the Macedonian Fake-News Complex," *Wired*, February 15, 2017, https://www.wired.com/2017/02/veles-macedonia-fake-news.

6. Jürgen Pfeffer, Thomas Zorbach, and Kathleen M. Carley, "Understanding Online Firestorms: Negative Word-of-Mouth Dynamics in Social Media Networks," *Journal of Marketing Communications* 20, no. 1–2 (March 4, 2014): 117–28, https://doi.org/10.1080/13527266.2013.797778.

7. Patrick M. Morgan, "Deterrence and Rationality," *Deterrence Now* (Cambridge, UK: Cambridge University Press, 2003), 42–79.

8. Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011).

9. Michael J. Mazarr, *Rethinking Risk in National Security: Lessons of the Financial Crisis for Risk Management* (Basingstoke, UK: Palgrave Macmillan, 2016).

10. Robert Jervis, "War and Misperception," *Journal of Interdisciplinary History* 18, no. 4 (Spring 1988): 675–700.

11. Sonya Song and Steven Wildman, "Using Online Media Audience Data to Develop and Refine Media Strategy," in *Media Business Models: Connecting Media to their Markets*, ed. C. Scholz and S. S. Wildman (Lisbon, Portugal: Media XXI, in press).

12. See, for example, Herbert Lin and Jaclyn Kerr, "On Cyber-Enabled Information Warfare and Information Operations," in *Oxford Handbook of Cybersecurity* (New York: Oxford University Press, forthcoming).

13. Samuel C. Woolley and Philip N. Howard, "Computational Propaganda Worldwide: Executive Summary," Computational Propaganda Research Project, Oxford Internet Institute, Oxford University, June 2017.

14. Heather Williams, "'Blind Moles' with Smartphones: Social Media and Nuclear Crisis Escalation," paper presented at Effects of the Global Information Ecosystem on the Risk of Nuclear Conflict conference sponsored by the Stanley Center for Peace and Security, the Center for International Security and Cooperation, and the Hoover Institution, Stanford University, September 7, 2018.

15. Lawrence Freedman, “Escalators and Quagmires: Expectations and the Use of Force,” *International Affairs* 67, no. 1 (January 1991): 15–31.

16. See chapter 7, this volume.

17. US State Department, “The Cuban Missile Crisis, October 1962,” Office of the Historian, https://history.state.gov/milestones/1961-1968/cuban-missile-crisis.

18. Jessica Chen Weiss and Allan Dafoe, “Authoritarian Audiences, Rhetoric, and Propaganda in International Crises: Evidence from China,” *International Studies Quarterly* 63, no. 4 (December 2019): 963–73.