

CHAPTER 8

In Screen We Trust

The world isn't run by weapons anymore, or energy, or money. It's run by ones and zeros—little bits of data. It's all electrons. There's a war out there, a world war. It's not about who has the most bullets. It's about who controls the information—what we see and hear, how we work, what we think. It's all about information.

AS STATED BY COSMO (BEN KINGSLEY),
THE VILLAIN IN *SNEAKERS*

All systems check. The five thousand spinning centrifuges in operation at Iran's nuclear enrichment facility at Natanz were humming away, and the Islamic Republic was making good progress on its "peaceful" nuclear energy program. If things continued on track, soon Iran would have enough enriched uranium 235 (U-235) to create its own nuclear power plant or its first atom bomb, depending on whom you asked. Although Iran had always insisted its nuclear activities were purely for civilian energy use, much of the world, including the United States, Europe, Israel, and the United Nations, was less convinced.

In 2005, the UN's International Atomic Energy Agency (IAEA) found Iran in noncompliance with the Nuclear Non-proliferation Treaty it had signed, and the inspection agency reported its concerns to the UN Security Council. In response, the UN demanded Iran

suspend its nuclear enrichment activities at Natanz, to which the nation's president at the time, Mahmoud Ahmadinejad, responded with an emphatic no. Senior officials at the IAEA concluded Iran had sufficient information to design and produce a workable atom bomb and UN sanctions were imposed. But would the sanctions prevent Iran from getting the bomb? Given Iran's preeminence on America's "Axis of Evil" hit list, something more had to be done.

For political reasons, an overt military strike was ruled out, but the following year President George W. Bush authorized a covert attack on the nuclear facilities at Natanz and dubbed the top secret program Operation Olympic Games, according to the *New York Times*. The result was the "most significant covert manipulation of the electromagnetic spectrum since World War II, when cryptanalysts broke the Enigma cipher that allowed access to Nazi codes."

The Iranians were no easy target and were smart enough to not connect the most prized information network in the Islamic Republic to the Internet. As a result, the operatives associated with Operation Olympic Games could not just tunnel their way in via a poorly protected road on the information superhighway. A network of human agents, engineers, and maintenance workers—spies and unwitting accomplices alike—would have to be assembled and choreographed with tremendous precision if the plan were to succeed. The weapon of choice for this covert operation? A USB thumb drive.

To sabotage the centrifuges at Natanz, a new class of cyber weapon was created, one that could leap from the virtual world of computers and enter the physical world of industrial control systems. Enter Stuxnet, a highly sophisticated computer worm widely believed to have been created by the United States and Israel to keep a notorious foe in check. The authors of Stuxnet copied the worm onto a simple USB flash drive, now locked and loaded, ready to seek out its quarry. How the drive came to be smuggled into Natanz and who inserted it into the computer network at the facility remain unknown, even today.

What is known, however, is how quickly the malware spread across the IT infrastructure of the plant. The mere insertion of the flash drive into a single computer's USB port infected the machine's Microsoft Windows operating system using a zero-day exploit not

previously documented. The worm also used a forged digital security certificate, indicating it was reliable and trustworthy, allowing it to replicate across the IT infrastructure at Natanz with impunity. As the worm propagated from desktop to desktop and network to network, it posed a simple question of each machine it infected: Is this computer connected to an industrial control system manufactured by the German multinational company Siemens?

The Americans and the Israelis had done their homework and knew the centrifuges at Natanz were run by Siemens S7-417 industrial programmable logic controllers (PLC), which monitored the valves and pressure sensors of the centrifuges at the plant. If a computer was not connected to a Siemens PLC, the worm failed to replicate and merely died. If, however, Stuxnet detected that a desktop or network computer system was connected to a Siemens PLC, the cyber weapon assiduously began its work, making its way from the Windows computer onto the industrial control system that managed the Iranian centrifuges.

The perpetrators of the attack knew that refining U-235 was a very tricky business. The IR-1 centrifuges used at Natanz were designed to spin at 100,000 rotations per minute (RPM), an incredible feat in speed and technology. If the centrifuges spin too slowly, the U-235 required for nuclear energy (and bombs) does not separate effectively. If they spin too fast, the centrifuges begin to vibrate and shake uncontrollably until the pressure becomes so severe the motors burn out, requiring the centrifuge to be replaced. The authors of Stuxnet understood that no centrifuges meant no enrichment, thus no bomb and no threat.

The Siemens PLCs were key to the attack, but the authors of Stuxnet were not impetuous cyber warriors with a pillage-and-burn mentality. They were patient, strategic, and cunning in their attack on Natanz. In the first phase of the assault on Natanz, Stuxnet did nothing but observe, sitting there silently, stealthily gathering information to understand how the enrichment centrifuges worked. The worm recorded all of its findings in a masterful preplanned move that would prove crucial to the success of the operation.

It was in phase two, however, that Stuxnet began to show its true powers as the worm established dominion over the industrial

control systems at Natanz. Slowly, its puppet masters began manipulating the centrifuge valves and motors responsible for enriching U-235 at the facility. For months, and even years, the centrifuges sped up and slowed down, fluctuating from their designed 100,000 RPM specifications. Centrifuge pressure mounted, rotors failed, and yields of enriched uranium began to plummet.

Meanwhile, inside the highly secure operations control room at Natanz, all systems were in full working order—at least according to the computer screens monitored by the engineers at the facility. Every one of the thousands of centrifuges was represented by a light on a computer screen, and each was carefully monitored for system malfunctions. A green light meant the centrifuge was working as designed; a gray or red light indicated problems. Day after day, engineers dutifully watched their screens for any evidence of failure, as the lights continued to shine bright green on the data safety systems before their eyes. Cascade protection system? Check. Centrifugal pressure? Check. Rotor speed? Check. Screens on the walls, screens on their desktops, screens on their control panels—all information systems inside their operations command center told the Iranians their nuclear ambitions were on track. Yet nothing could be further from the truth.

The damage caused by the Stuxnet worm was designed to be subdued at first. Gradually, some centrifuges began spinning out of control, but the Iranians blamed bad parts or the incompetence of their engineers. Each centrifuge that failed seemed to have a different explanation: this device was too slow, that one was too fast, there was too much pressure in these. The uranium processed was increasingly of poor quality and not usable. Inspection after inspection of the facility was carried out, and researchers continued to closely observe the status of their entire operation from the computers inside their control room. As time passed, dozens and then hundreds of centrifuges began to fail. Iran's nuclear ambitions were now in doubt. What the hell was going on? As it turned out, the Iranians had placed too much trust in the computer screens governing their prized secretive nuclear enrichment site.

The data logging and computer recording of the industrial control systems stealthily perpetrated by the Stuxnet worm in phase one of the attack had a clear, if not immediately obvious, pur-

pose: to fully document what the Siemens PLCs looked like when they were in full, proper working order. Rotors spinning according to plan and pressure at expected levels yielded all systems go, all maintenance lights green. Stuxnet captured all of those data and recorded it on the PLC equivalent of a VCR, carefully saved for posterity. What happened next was straight out of a Hollywood blockbuster, portrayed many times in films such as *Ocean's Eleven* and *National Treasure*. The attackers simply prerecorded video footage of the casino vault or safe room to be targeted and played it back on the screens of the watchers and security staff.

As the uranium enrichment centrifuges spun out of control at Natanz, Stuxnet masterfully intercepted the actual input values from the pressure, rotational, and vibration sensors before they reached the operational control room monitored by the plant's engineers. Rather than presenting the correct real-time data from the Siemens PLCs, Stuxnet merely replayed the prerecorded information it had taken during phase one of the operation, showing all systems in full working order. The brilliant move meant that even though in reality the industrial control systems were melting down and digitally screaming for help, the flashing red danger signs displayed by the system were supplanted by a sea of green calm on the monitors of the Iranians controlling Natanz. As the centrifuges spun out of control and tore themselves apart, the human operators in the digital control room had no idea their own reality had been hacked, hijacked by a computer worm with a funny name sent on a mission to search and destroy.

Life in a Mediated World

Unfortunately, you have more in common with the Iranians than you realize. While you may not be producing U-235, you too depend on screens every day to translate the world around you. Your cell phone tells you who has called, your PC reminds you that you need to update your operating system, and the GPS in your car shows you how to get to your morning meeting. All of this and more transpire long before you finish your second cup of coffee. The result? We no longer live life through our own innate primary

human sensory abilities. Rather, we experience it mediated through screens, virtual walls that sever us from our intrinsic senses and define the world for us. Screens interpose themselves between us and the real world, projecting information that is purportedly equal to reality but is at best only ever a rough approximation, one that is easily manipulated.

At our airports, hospitals, banks, and ATMs, screens have become an omnipresent fixture in our lives. But screens today are dumb; they do little more than present the underlying information contained in data systems, systems which are eminently hackable. Those who control the computer code also control our screens and thereby our experiences and our perceptions. Everything from video games to voting machines can be tampered with, and in this brave new world seeing something with your own two eyes and hearing it with your own ears is by no means an indication that it is legitimate, correct, or safe. As a result, the screens we watch can deceive us in ways most have yet to understand.

Whether or not you realize it, your entire experience in the online world and displayed on digital screens is being curated for you. Some of this filtering of course is good. With billions of tweets, Snapchats, status updates, and blog posts, there is no way any of us could consume the volume of data thrown our way on a daily basis. Knowing this, Internet companies go to great lengths to learn what you like and to customize your online experience using a series of computer algorithms. Facebook studies your Web links, images, pokes, messages, events, and Likes to customize what you see on your screen every day. As a result, you do not see most of what's posted by your friends or on the pages you follow, and your friends see perhaps only 10 percent of your own updates on their news feeds. For as much effort as Facebook puts into studying and segmenting you for its advertisers, it works at least as hard to determine which of your friend's posts you would most likely want to see every time you visit his site or launch his app. But why does it do this? Simply stated, Facebook, Google, and other Internet companies know that if they provide you the "right" stuff, you'll spend more time on their sites and click on more links, allowing them to serve you up more ads.

Facebook is by no means alone in this game, and Google too

quantifies all your prior searches and, more important, what you've clicked on, in order to customize your online experience. In his book *The Filter Bubble*, the technology researcher Eli Pariser carefully documented the phenomenon. Getting you the "right" results is big business, and millions of computer algorithms are dedicated to the task. Google reportedly has at least fifty-seven separate personalization signals it tracks and considers before answering your questions, potentially to include the type of computer you are on, the browser you are using, the time of day, the resolution of your computer monitor, messages received in Gmail, videos watched on YouTube, and your physical location. Google alters, in real time, the search results it provides you based on what it knows about you. A search for the word "abortion" returns links to Planned Parenthood for some and Catholic.com for others; if your query is "Egypt," you may receive results on the Arab Spring, while your mom sees info on the pyramids or Nile cruises. Like Pariser, you can run this experiment yourself, and the results will provide an illuminating perspective on how Google sees you.

The fact of the matter is that there is no such thing as "standard Google." Eric Schmidt has publicly acknowledged that "it will be very hard for people to watch or consume something [online] that has not in some sense been tailored for them." While none of this is necessarily malicious, there are important questions to be asked about how this information is being culled, sorted, and curated by others purportedly on your behalf. The challenge, however, is that Google, Facebook, Netflix, and Amazon do not publish their algorithms. In fact, the methods they use to filter the information you see are deeply proprietary and the "secret sauce" that drives each company's profitability. The problem with this invisible "black box" algorithmic approach to information is that we do not know what has been edited out for us and what we are *not* seeing. As a result, our digital lives, mediated through a sea of screens, are being actively manipulated and filtered on a daily basis in ways that are both opaque and indecipherable. This fundamental shift in the way information flows online shapes not only the way we are informed but the way we view the world. Most of us are living in filter bubbles, and we don't even realize it.

Around the world, nations are increasingly deciding what

data citizens should be able to access and what information should be prohibited. Using compelling arguments such as "protecting national security," "ensuring intellectual property rights," "protecting religious values," and the perennial favorite, "protecting children," governments are ever expanding their national firewalls for the purpose of Internet censorship. Some of these filtering techniques are disclosed to the general public. For example, in France and Germany, sites promoting Nazism or denying the Holocaust are openly censored. In Syria, YouTube, Facebook, Amazon, Hotmail, and pro-Kurdish sites have been blocked. In Saudi Arabia, 400,000 sites have been restricted, including those that discuss any political, religious, or social issue incompatible with Islam or the personal beliefs of the monarch. In many instances, however, there is no indication that your online information is being censored; instead, your content simply does not appear. In the United Arab Emirates, the government has even blocked all access to the entire .il domain of Israel, digitally erasing the existence of the Jewish state from the virtual world.

Tech companies have collaborated in national censorship programs and acceded to state demands to filter offending content in real time, as Google did when entering the Chinese market in 2005. But perhaps no other government is as adept at and rigorous with its Internet-filtering programs as China. The "Great Firewall" of China ensures that its billion-plus residents are unable to see politically sensitive topics, such as the Tiananmen Square protests, embarrassing details about the Chinese leadership, or discussions of Tibetan rights, the Dalai Lama, Falun Gong, Taiwanese independence, political reform, or human rights. Internet censorship, however, is not restricted to autocratic regimes or despots. As of 2014, there were more than four billion people living in countries that practice Internet filtering of one sort or another.

Screens tell you not what is really out there but what the government or Facebook thinks you should see. If you searched for something and it wasn't there, how would you know it really was? To paraphrase an old philosophical question, if a tree falls on the Internet and no search engine indexes it, does it make any noise? As we live our lives increasingly mediated through screens, when it doesn't exist online, it doesn't exist. If an event is not listed in

Google, it never happened. Conversely, if it does appear in Google, it still might not have happened. Welcome to the world of digital trickery, a virtual hall of mirrors represented as screens where all is magically possible.

The profound risk of life in a technologically mediated world is that it creates mammoth opportunities for information to be manipulated in undetectable ways that most neither expect nor understand. Screens are everywhere, beeping, ringing, and blinking for our attention. But what if these screens were lying? Feeding us false information and misleading us? In today's world, all that we see on screens can be faked and is easily spoofed. Ask anybody who has ever visited an online dating site, and he or she will tell you: what you see is not always what you get.

Does Not Compute

Why, sometimes I've believed as many as six impossible things before breakfast.

LEWIS CARROLL, *THROUGH THE LOOKING GLASS*

What do hackers, fraudsters, and organized criminals have in common with Facebook, Google, and the NSA? Each is perfectly capable of mediating and controlling the information you see on your computer screens. In a world where information is power, the gatekeepers who control the flow of data to your screen can also control others. We encounter this behavior on a daily basis every time we go online. Most of us would not consider making a major purchase or reserving a table at a new restaurant for a special occasion without first doing our own Internet research. Who better to inform us than our fellow shoppers and diners? Nearly 90 percent of consumers say online reviews influence their buying decisions, and a Nielsen study found that a surprising 70 percent trust the reviews they read online as much as recommendations from a friend. Unfortunately, according to an investigation by the New York State attorney general, 25 percent of the reviews on Yelp, one of the most popular sites of this kind, are completely bogus. Worse, in September 2014, a federal appeals court ruled it was completely

legal for Yelp to manipulate its ratings based on which companies advertised on the site; big spenders could legally get five stars, even if all users rated them a one. Reviews on eBay, Amazon, and TripAdvisor are also all easily faked, and many of those five-star postings you see were written by the businesses themselves or by paid proxies. There are even professional companies whose entire business model rests on gaming the online review system. The practice is known as astroturfing and is widespread. One company investigated by the State of New York, known as Zamdel Inc., was accused of writing more than fifteen hundred fake reviews on Yelp and Google Places.

I Thought You Were My Friend

According to Facebook's own 2014 annual report, up to 11.2 percent of its accounts are fake. Considering the world's largest social media company has 1.3 billion users, that means up to 140 million Facebook accounts are fraudulent and these users simply don't exist. With 140 million inhabitants, fake Facebook-land would be the tenth-largest country in the world. Just as Nielsen ratings on television sets determine different advertising rates for *The Walking Dead* versus the Super Bowl, online ad sales are determined by how many eyeballs a Web site or social media service can command—if only the data could be believed.

Want 4,000 followers on Twitter? They can be yours for \$5. Want 100,000 fans on Facebook? No problem, you can buy them on SocialMediaCorp.org for a mere \$1,500. Have even more cash to burn? How about a million new friends on Instagram? "For you we make special deal," only \$3,700. Whether you want favorites, Likes, retweets, up votes, or page views, all are for sale on Web sites like Swenzy, Fiverr, and Craigslist. These fraudulent social media accounts are then used to falsely endorse a product, service, or company, for a small fee of course. Most of the work is carried out in the developing world, in places such as India and Bangladesh, where actual humans may control the accounts. In other locales, such as Russia, Ukraine, and Romania, the entire process has been scripted by computer bots, little programs that will carry out your

pre-encoded automated instructions, such as "click the Like button," over and over again using different fake personas.

Just as mythological shape-shifters were able to physically transform themselves from one being into another, these modern screen shifters have their own magical powers, and criminals are eager for a taste of the action, studying their techniques and deploying them against easy marks for massive profit. In fact, many of these clicks are done for the purposes of "click fraud." Businesses pay companies such as Facebook and Google every time a potential customer clicks on one of those banner ads or links you see online, but organized crime groups have figured out how to game the system to drive profits their way via so-called ad networks, which capitalize on all those extra clicks. Stung by the criticism, social media companies have attempted to cut back on the number of fake profiles out there. The results of Facebook's actions were revealing. Rihanna and Shakira lost 22,000 Facebook fans, Lady Gaga had 32,000 of hers removed, and Zynga's *Texas Hold 'Em Poker* had 100,000 purported supporters vanish into thin air.

If Facebook has 140 million fake profiles, there is no way they could have been created manually one by one; there has to be something much more sinister at work, and there is. The practice is called sock puppetry and is a reference to the children's toy puppet created when a hand is inserted into a sock to bring the object to life. In the online world, organized crime groups create sock puppets by combining computer scripting, Web automation, and social networks to create legions of online personas. This can be done easily and cheaply enough to allow those with deceptive intentions to create hundreds of thousands of fake online citizens.

One only needs to consult a readily available online directory of the most common names in any particular country or region. Have your scripted bot merely pick a first name and a last name, then choose a date of birth and let the bot sign up for a free e-mail account. Next, scrape online photo sites such as Picasa, Instagram, Facebook, Google, and Flickr to choose an age-appropriate image to represent your new sock puppet. Armed with an e-mail address, name, date of birth, and photograph, you just need to sign up for an account on Facebook, Twitter, or Instagram. As a final step, you teach your puppets how to talk by scripting them to reach out and

send friend requests, repost other people's tweets, and randomly "like" things they see online. Your bots can even communicate and cross post with one another. Before you know it, you have thousands of sock puppets at your disposal for use as you see fit. It is these armies of sock puppets that criminals use as key constituents in their phishing attacks, to fake online reviews, to trick users into downloading spyware, and to commit a wide variety of financial frauds—all based upon misplaced trust.

Fatal System Error

We now live in an "in screen we trust" world. We look first and foremost to computers for guidance and direction. We depend on screens to give us answers, and all too infrequently do we question the results. But if your programming is poor or your primary data incorrect, these errors will be reflected in the results you receive. Garbage in, garbage out is one of the axioms of computer science. In the past, our limited reliance on technology insulated and protected us from many of these mistakes. In the age of big data, however, the calculus has shifted—big-time. We are all affected by database errors in one way or another, and the implications of these inaccuracies are growing every day. According to the Federal Trade Commission, nearly 25 percent of all consumer credit reports contain errors, and data brokers such as Acxiom have admitted that 30 percent of the data they maintain on you may be inaccurate.

When the forty million to fifty million Americans affected by these errors attempt to rent an apartment, buy a car, get a mortgage, or apply for a job, they soon discover that somebody else's mistake has now become their nightmare. If "according to our computer" you are a credit risk, you are, no ifs, ands, or buts. Today, millions of decisions are made every day with faulty, incomplete, or imprecise data, often with no further verification. If the problem were limited to credit reports, it might be tolerable, if only barely. But living in the land of "in screen we trust" means computer errors can affect not only your finances but also your life and your liberty.

As the world of medicine rushes to digitize patient records in an effort to save money, improve efficiency, and yield new big-data

insights on disease, there has been an unintended cost: accuracy. Tens of millions of electronic medical records contain incorrect information about patients, and erroneous data on computer screens can quite literally kill. Gary Foster, a twenty-seven-year-old from Essex, England, died at University College Hospital in London when a glitch in the facility's computer system meant the young man received an overdose of cancer drugs during his stay. The hospital staff, following an inaccurately entered prescription order, provided him with lethal doses of chemotherapy for his testicular cancer treatments. Not only can too much faith in computer screens kill, but also it can have a deleterious impact on public safety.

In California, a computer glitch led to the release of 450 dangerous criminals after a system error directed prison guards to set some of the state's most violent offenders free. Gang members, rapists, armed robbers, and inmates classified as having a "high risk for violence" walked out of prisons statewide because officials accepted the information on their screens as the truth. Of course, errors in criminal justice data are common, and these mistakes not only free the guilty but also inculcate the innocent. In Britain, police officials at the national Criminal Records Bureau admitted that more than twenty thousand people had been wrongly branded as criminals because of data errors in their system. The massive blunder meant that thousands of innocent individuals were given criminal records for offenses they never committed. "But, Officer, you've got the wrong person" is a refrain cops are used to hearing; unfortunately for those involved, what is on the screen is the truth, until proven otherwise. Across the U.K., victims of these errors were denied jobs and volunteer positions and had their reputations destroyed, all because of our unyielding faith in our screens.

Today we face a confluence of phenomena, both human and technical, that are coming together like a perfect storm to pose particularly striking dangers to our society. With each successive generation, we grow deeply comfortable, even if only unconsciously, with blindly following the directions provided to us by machines. Garbage in, garbage out has been supplanted by garbage in, gospel out: if the computer says so, it must be true. The problem with such reasoning is that we as a society are relying on incorrect data all

the time, a festering problem that will come back to bite us. Filter bubbles, invisible search engine censorship, national firewalls, and faulty data mean we have a fundamental integrity problem with the way we see the world, or more precisely with the way the world is presented to us, mediated through our screens.

When Seeing Ain't Believing

In the preceding chapters, we focused extensively on what happens when your data leak and your information confidentiality is breached. No doubt, criminals are having a field day with all the opportunities they have generated by stealing your data. But there is a much more profound and insidious threat to the world's information—changing it. Criminals, hackers, terrorists, and governments are increasingly breaking into data systems, not to steal information, but to surreptitiously manipulate its presentation on our screens, as we saw in Natanz. As a result, the very integrity of the world's information is under attack. Slowly, imperceptibly, and with great precision, attackers can enter our data systems and covertly modify any and all of the underlying information. When hackers attack, stealing our data may be the best-case scenario compared with altering the information without our knowledge.

In the 1995 movie *The Net*, Sandra Bullock plays a reclusive systems analyst who accidentally uncovers a plot by a diabolical cyber-terrorist organization to take over the world's information systems. The film opens with the undersecretary of defense committing suicide after learning he has tested positive for HIV at the Bethesda Naval Hospital. As it turns out, the official was never HIV positive; rather, hackers changed the results on his medical tests, in retaliation for the undersecretary's pursuit of international cyber villains, information that his doctor dutifully reported based on the data on his computer screen. The shame of the test results was too much for the conservative undersecretary, precipitating his suicide.

This is the world of information warfare, where computer disinformation disseminated through an array of blinking screens carries real-world impact. The events depicted in the film are

decidedly possible today. Police data systems have been successfully hacked globally, including in Australia, England, Italy, Memphis, Montreal, Hong Kong, and Honolulu. In 2013, the Danish police national driver's license registry was breached, and it was believed the hackers made changes to the underlying law enforcement data systems. In Philadelphia, also in 2013, a classified database of witnesses to some of the city's most notorious crimes was breached by a local criminal group. As a result, the names, addresses, and photographs of dozens of protected witnesses were posted on Instagram with the tagline "Expose the Rats." Many of the individuals exposed had testified in secret grand jury hearings, and within a few days there were nearly eight thousand followers for the Instagram account user known as rats215. One nineteen-year-old witness who had testified in a homicide case was later targeted for retaliation in a shooting. In what amounted to mass witness intimidation, numerous visitors to the site posted comments such as "exterminate the rats" and "put out a hit on them."

In Massachusetts, a prisoner already serving time for computer hacking was allowed access to the prison's library computer for the purposes of legal research on his own case. Once his fingers touched the keyboard, he was able to tunnel through the computer network at the Department of Correction and obtain access to case files on other inmates, as well as the names, dates of birth, Social Security numbers, home addresses, and phone numbers of the prison's eleven hundred guards. Given the insecurity of criminal justice systems, how many prisoners have been wrongly released, like the 450 violent offenders in California, because the underlying data were falsified and purposefully tampered with? The answer is, we simply do not know, and government officials are loath to discuss the matter.

As open and vulnerable as law enforcement computers are, they are a veritable Fort Knox compared with our electronic medical records. Forget for the moment the millions of accidental errors previously noted; the Department of Health and Human Services (HHS) has determined that at least twenty-one million Americans have had their electronic medical records accessed without authorization since 2009. In fact, HHS has documented more than nine hundred such breaches at hospitals across the United States.

But how many more weren't reported? Federal law only mandates reporting if more than five hundred records per incident are targeted. Organized criminals are targeting medical data in a wide variety of ways, ranging from Medicare fraud to extortion. In Virginia, hackers accessed eight million patient records and thirty-five million prescriptions maintained by the state's Department of Health and threatened to post the information online unless Virginia paid a \$10 million ransom. Globally, electronic medical data systems are deeply porous, and bad actors are fully capable of leveraging these data with deadly consequences.

Time and time again, doctors, nurses, and technicians will follow the directions presented to them on computer screens, even when the information is incorrect, as we saw previously when describing the fatal hospital system errors that led to the death of Gary Foster. If the screen says you are HIV positive, the hospital will deliver that news to you. Worse, if your blood type is listed as O positive and a hacker, enemy, or adversary switches it in the hospital's database to A negative before you go into surgery, the operation will likely result in death. The same would be true if somebody maliciously erased your allergy to penicillin from your digital chart and a nurse innocuously carried out a medical order directing her to inject five hundred milligrams of the drug into your IV.

The profound consequences of the "in screen we trust" mentality can open the door to an array of new crimes, including new ways to commit murder. In response, criminals have developed a panoply of methodologies to profit from a world that has subsumed human intelligence in favor of the digital and the virtual. Nefarious actors are proving particularly adept at so-called man-in-the-middle attacks, wherein they insert themselves between reality and the data we see on our screens. The result? An all-out assault on the integrity of the information we're stockpiling as a result of the big-data revolution.

Screen of the Crime

For every screen in your life, criminals have developed a plan of attack. One of the most common such scams on the Internet is the

phenomenon of phishing—a technique by which criminals masquerade as a legitimate Web site in order to acquire information such as passwords and credit card numbers. The term “phishing” is a hacker spelling of the word “fishing,” and the technique involves trying to get an innocent fish to take the bait of a malicious link and bite. The organized crime groups that run phishing cons try to trick users into clicking on a link that takes them to a fake Web site controlled by fraudsters. Phishing messages arrive in our in-boxes, via SMS, tweets, instant messages, and Facebook status updates. They allegedly come from our banks, cable companies, retirement plans, social media outlets, and mobile phone operators and target users around the world, with the greatest number of victims in the United States, the U.K., and Germany.

In the end, all phishing attacks depend on an unsuspecting user clicking on a link or attachment in a message that will either take the unsuspecting party to a fraudulent Web site or install malware on the user's machine. Criminals take advantage of HTML hypertext links and embed their attacks in hidden computer code. Phishing messages arrive as fake e-cards, e-mails from our bank, job offers, coupons, or deals too good to be true on social media. These malicious communiqués, replete with grammatical and spelling errors in years past, have become highly professionalized and are today virtually indistinguishable from the real thing. Criminals know exactly how to subvert the trust you have placed in your screens by visually mimicking the sites they impersonate and tricking your senses with a digital sleight of hand.

A typical message might arrive from an address such as security@bankofamerica.com, informing you that you need to update your profile or that your account has been suspended because of suspicious activity. Uh-oh, seems important, better look into it, you think. What you may not know is that the e-mail address that shows up in your in-box is ridiculously easy to spoof or fake. Anytime you set up a new mail account in any mail software program such as Outlook, Mac Mail, or Thunderbird, you will be asked to enter a name and e-mail address. If a crook types “Bank of America Security Team” as the name in an e-mail program, that's what will show up in your in-box. It's just that simple. Only by examining the message headers might you note that the e-mail address used

by the bad guys was actually notifications@security-bankofamerica.com—still close enough to fool the average bear.

In every way, the message looks as if it were from your bank—same font, same color, same logo—but it's not. Though the visible link may say www.bankofamerica.com, you will instead be taken to www.bank0famerica.com (0 instead of o) or even bankofamerica.accountupdates.com (accountupdates.com is the actual domain you are visiting, owned by the criminals; Bank of America is just a folder on their site to fool you); www.citibank.com will be supplanted by www.citiibank.com (two i's in the fake address, barely noticeable). Phishing messages all subtly scream what they want you to do by making the embedded link containing the poison pill impossible to miss, written in a large font or having a big colorful button: "To update security settings and protect your account, click here." And now they own you.

The fateful click will take you to the Citiibank.com Web site, where you will be prompted to log in and provide your credentials, and when you do, the thieves capture your log-in name and password, as well as a bevy of other personal information. And now is when the criminals really go to work. Phishing is a gateway crime, a fundamental first step that provides thieves with the data they need to perpetrate step two in their plot against you, including identity theft, financial fraud, tax fraud, and insurance fraud. Just as the nuclear engineers in Natanz were presented with a compelling but completely fictional reality on their screens, you too are beset by criminals knocking at your door every day using similar deception techniques.

For criminals, the costs of pulling off such digital charades are ridiculously low. Fully automated phishing kits sold on the digital underground to send scam messages to 500,000 e-mail addresses cost a mere \$65, and as mentioned before, criminals take advantage of "sock puppet" accounts to scale up. As a result, more than 100 million phishing messages arrive in our in-boxes every day. According to a study on the economics of these attacks by Cisco, approximately eight people out of a million will succumb to the ruse, with an average loss of \$2,000 per victim. Thus for about \$130, crooks can generate \$16,000, a 12,000 percent return on investment. With 36 billion phishing messages sent annually, the scale, scope, and

profitability of cyber crime come into focus. Although the returns from bulk phishing attacks are impressive, they pale in comparison to "spear phishing," a technique that does not send bulk fraudulent messages to millions, but rather carefully targets specific individuals or organizations.

Spear-phishing attacks have become the go-to tool of choice for those committing industrial espionage, and the costs in those cases can be gargantuan, as the global beverage giant Coca-Cola discovered. As part of its Asian expansion, Coke had entered into advanced-stage negotiations to purchase China's Huiyuan Juice Group. All was proceeding according to plan with the acquisition, until it inexplicably fell apart. Something fishy was going on, and Coke wanted answers. It conducted a full investigation into the matter and began to examine the deal in detail, including communications between Coke and representatives at the Huiyuan Juice Group. In the end, Coke discovered that the Chinese government was aggressively monitoring the deal and pursuing nonpublic insights into Coke's bidding plans and intentions. How did the Chinese ultimately get the access they needed? By manipulating the screen of Paul Etchells, the deputy president of Coca-Cola's Pacific Group.

Etchells opened an e-mail that had been spoofed to appear as if it had come from a senior executive in Coke's legal department. The subject line was enticing: "Save Power, Save Money—from Coke's CEO." Etchells knew his boss at Coke had been heavily pushing energy savings in the company (as did the Chinese who had infiltrated Coke's corporate information systems). The perpetrators of the attack spoofed reality by making the message appear as if it came from a trusted colleague, on the internal corporate network, with a subject line that was compelling and contextually made sense. When Coke's deputy president innocently clicked on the link, he also silently downloaded malware onto his workstation, including a keystroke logger that captured everything the executive typed. As a result, the Chinese were able to download a barrage of computer files related to the deal. While Coke has refused to publicly comment on "security matters," what is clear is that one single spear-phishing attack against a senior Coke official cost the firm its \$2.4 billion acquisition of China's Huiyuan Juice Group.

Coke is not alone in having succumbed to spear phishing: it has become the preferred method of attack for online criminals and digital spies, responsible for a full 91 percent of all targeted cyber attacks.

Criminals can now even change what you see on your screen in real time, including your financial statements. But what if your bank account balance were zero, and you didn't know it? There are so many thousands of malware programs that will steal money from your bank account these days; the whole process has become both routine and even automated by the bad guys. Criminals infect your computer or mobile phone, capture your log-in credentials, and then use them to drain your account balance. Of course should you yourself happen to log in and notice the low balance, you may have a shot at notifying your bank's fraud department and stopping the flow of the funds. Banks usually reserve settlement periods of a day or so, particularly internationally, in which they can cancel, stop, or reverse a transfer or wire, but the time frame to do so is incredibly short. To this end, criminals will go to great lengths to ensure that what you see on your computer screen is not what you have in your bank account. Highly specialized Trojan horse software, such as SpyEye and URLZone, not only steals your money but will even offer you false reassurances that it is still in your bank account. The magic of these Trojans is that they give thieves more time to use your banking, debit, and credit card information without your realizing what is happening. Your first clue there is a problem is when you try to use your own bank card to make a withdrawal, only to learn you are overdrawn and there are insufficient funds.

The crimeware (criminal software) involved is so sophisticated it even knows how much it has stolen from each bank account compromised. Thus, if the thieves stole \$2,419 from your checking account, an algorithm will add that portion back to what you see on your screen in real time as you view your online account balance. Purchases made by criminals with your credit or debit card are automatically struck from the recent transactions list and the online statement before they appear on your screen. Even PDF copies of your banking and credit card transactions sent to your printer are modified before they come out of your machine. When these thieves own you, they really own you.

These types of man-in-the-middle attacks are powerful reminders that criminal hackers are perfectly capable of intermeddling reality for you via the ever-increasing number of screens in your life. Just like the perpetrators of Stuxnet, these criminals recognize that screens are merely a proxy for reality, one that is completely malleable and easily manipulated. Yet not all manipulation of the data we see on our screens is carried out by global cyber-crime cartels or espionage services.

Pedophiles routinely take on the digital personas of children, and 80 percent of the time those under eighteen cannot tell when they are speaking to an adult, making the screens of the young particularly vulnerable to attack. Recall the case of Amanda Todd, the twelve-year-old girl who was tricked into showing her breasts on video camera to a person she thought was a boy her age. Todd was blackmailed and tormented by her virtual attacker, driving the Canadian student to take her own life. For years, the case went unsolved, and Amanda's parents never knew who had tortured their daughter, until April 2014, when a break in the case led the Royal Canadian Mounted Police to the Netherlands and a suspect five thousand miles away. Police in Holland identified thirty-five-year-old Aydin Coban and charged him with "one count each of extortion, internet luring, criminal harassment, possession of child pornography for the purpose of distribution." Coban's alleged modus operandi was to establish a fake online persona, gain the trust of underage girls, and then seduce them into performing sex acts in front of a Webcam. Dozens of other victims both in Canada and around the world are believed to have been targeted by the Dutch pedophile.

Even for adults, interpersonal relationships and screen manipulation can be a dangerous mix. Such was the case with Elizabeth Thrasher, who was accused of lashing out at the daughter of her ex-husband's new girlfriend. The jealous woman allegedly copied two photographs from the teen's Myspace account and posted them to Craigslist in the Casual Encounters section of the site. She also listed the innocent girl's home address, phone number, e-mail, and work information, stating that she was looking for sex. Thrasher's prey had no idea about the Craigslist posting until she started to receive phone calls, text messages, and photographs (including

nude pictures), along with solicitations for sex. In court, the girl testified that she felt as if she'd been "set up to get killed and raped by somebody."

Stock Screeners

Not only can individuals and companies be manipulated based upon what appears on their screens, but so too can financial markets. While whispers and suppositions have moved markets in the past, the breakneck speed of the Internet means the world often reacts before information has been verified. In August 2000, a hacker by the name of Mark S. Jakob, a twenty-three-year-old community college student from El Segundo, California, created a fake press release and forwarded it to the Internet Wire, a Web-based distributor of corporate announcements. Jakob chose the Emulex Corporation, a Nasdaq-traded communications equipment manufacturer, as his target. The hacker merely copied the stationery and style of previous Emulex press releases, spoofed the company's e-mail address, and forwarded his news story to the Internet Wire. The fictional press release stated that the Securities and Exchange Commission (SEC) had opened an investigation into Emulex, that its quarterly earnings were to be restated, and that the company's CEO, Paul Folino, had resigned in response. The sensational story went viral and was picked up by numerous other newswires, including TheStreet.com, CNBC, Bloomberg, and the Dow Jones Newswires.

The market's response was both predictable and immediate. "In a sixteen-minute period following the republication of the fake press release, 2.3 million shares of Emulex stock were traded, and the price plummeted \$61, from \$104 to \$43, resulting in Emulex's losing \$2.2 billion in market capitalization." Exactly the response Jakob had been hoping for when he shorted the stock, earning the young market manipulator a paper profit of \$250,000. Immediately, Emulex's CEO appeared on Bloomberg and in other financial press denying the story, but by then the damage had been done. Within six days, the FBI, working with the SEC, had identified Jakob, who was arrested and pleaded guilty to wire and securities fraud. When

all was said and done, legitimate investors in the market lost more than \$110 million because a kid at a community college manipulated the trust they had in their screens.

Screen manipulation in the financial services sector is commonplace, and so-called pump-and-dump schemes are the meat and potatoes of online securities fraud. The practice involves traders artificially pumping up the price of a stock through false and misleading positive statements posted online and then dumping their overvalued shares before their lies are discovered. The practice has flourished in cyberspace, and the FBI has arrested dozens of criminals for taking part in these scams. Though pump and dump is generally unsophisticated in its approach, both individuals and organized crime groups have earned hundreds of millions of dollars by manipulating the information we all see online.

Sometimes financial screens can manipulate you in ways you don't even realize, by watching you while you watch them. That's what professional traders at Goldman Sachs and JPMorgan learned about the Bloomberg trading terminals they had been using for years. Bloomberg terminals are the lifeblood of Wall Street, and firms pay \$20,000 per year per terminal to mine the reams of data they provide in order to make their daily trades. What these traders didn't realize, however, was that reporters from Bloomberg's news division had been granted administrative access that allowed them to monitor client activities as traders used their Bloomberg boxes. In other words, Bloomberg news reporters were keeping tabs on terminal use to facilitate news reporting. Traders who thought they were viewing information privately on a dumb terminal learned the terminal wasn't quite so dumb and was actually watching back.

More than 300,000 of the most influential people in the financial world, including bankers, hedge fund managers, and Treasury Department officials, rely on these Bloomberg boxes to carry out their deeply proprietary research, with each query tied to a specific individual. The scandal came to light when a Bloomberg journalist phoned Goldman Sachs to inquire whether a partner still worked there, noting he hadn't logged in to his terminal for days. The casual observation set off alarm bells at Goldman Sachs, which went public with the story.

It was later revealed that Bloomberg's twenty-four hundred

reporters were able to see histories of user log-in information on the company's terminals, as well as the various search functions they used, such as for equities or commodities. Goldman officials complained that Bloomberg reporters eavesdropped on customers using their terminals and used this private information to spy on the activity of individual Goldman partners, information they used to generate Bloomberg news stories. One former Bloomberg reporter noted that "there was always a discussion in the newsroom of how to use the terminals to break news."

Financial screens can also be hacked and manipulated via high-frequency trading (HFT). In his seminal 2014 book, *Flash Boys*, Michael Lewis tells the story of how Wall Street insiders have gamed the entire financial trading system by hacking time. By spending hundreds of millions of dollars on vastly superior technical infrastructure, high-frequency traders were able to shave mere milliseconds off their trading times, providing them with an exploitable advantage over their peers. *Flash Boys* follows Brad Katsuyama, a trader at the Royal Bank of Canada's New York office, and his incredibly complex multiyear investigation into the world of HFT. What he discovered was startling: the stock market as it appeared on the screens on his desk was an illusion.

As it turned out, every time Katsuyama attempted to execute a trade, the stock price moved before his order was completed. How did this happen? High-speed traders had figured out a way to exploit the variable speeds at which trading information moved along fiber-optic cables to the stock exchanges. Though the signals traveled at two-thirds the speed of light, over longer distances minuscule time lags added up and could be taken advantage of. By paying huge sums of money for the fastest cables, the most powerful computers, and the privilege of colocating their data servers on-site with the exchanges themselves, high-speed traders were able to see Katsuyama's intent to purchase a stock at a particular strike price and purchase it out from underneath him at the price displayed on his screen. Katsuyama was not alone; we were all being affected by the same problem; he was just the first to document it. High-frequency traders were front running the market and in return screwing us all, including you and your mutual fund purchases, your 401(k) plan, and even city pension plans.

High-frequency traders had hacked time and screen, in what amounted to a man-in-the-middle attack. They inserted themselves between the purported real-time stock market data projected on Katsuyama's screens and a much faster reality, which they controlled and owned. Their computers were so fast that they could detect other people's orders, cut in line before them, buy the stocks in question, and sell them back to the person who had originally attempted the purchase, albeit at a higher price. A few pennies' difference here and there on millions of trades a day meant high-frequency traders made billions of dollars in cumulative profits out of a five-millisecond trading advantage. To put that level of speed into context, the blink of a human eye takes approximately three hundred to four hundred milliseconds. It's akin to the scene in the movie *The Matrix* when the bad guys start shooting at Neo (Keanu Reeves) and he can see the bullets coming and can move at light speed to avoid them, except in this case it's all about hacking the financial system and none of us little people have Neo's powers.

In the days after *Flash Boys* was released, a series of investigations were launched by the SEC, the FBI, and the New York State attorney general. But their sudden interest begs an important question: How is it this entire system was even able to develop in plain view of the SEC right on the heels of the 2008 global financial crisis? Michael Lewis rightfully pointed out that the "market is rigged," but before this could happen, the firms involved in flash trading had to undermine your screen in order to create the fiction of a transparent and trustworthy marketplace. It is troubling to know that we live in a world in which the screens of hospitals, prisons, police departments, banks, brokerages, and news sites are so readily hackable, but as we shall see, screens are proliferating, threats are mounting, and these attacks can cost us much more than just our money.