1 Terrorism and Security

What the Terrorists Want

Originally published in Wired, *24 August 2006*

On August 16, two men were escorted off a plane headed for Manchester, England, because some passengers thought they looked either Asian or Middle Eastern, might have been talking Arabic, wore leather jackets, and looked at their watches—and the passengers refused to fly with them on board. The men were questioned for several hours and then released.

On August 15, an entire airport terminal was evacuated because someone's cosmetics triggered a false positive for explosives. The same day, a Muslim man was removed from an airplane in Denver for reciting prayers. The Transportation Security Administration decided that the flight crew overreacted, but he still had to spend the night in Denver before flying home the next day. The next day, a Port of Seattle terminal was evacuated because a couple of dogs gave a false alarm for explosives.

On August 19, a plane made an emergency landing in Tampa, Florida, after the crew became suspicious because two of the lavatory doors were locked. The plane was searched, but nothing was found. Meanwhile, a man who tampered with a bathroom smoke detector on a flight to San Antonio was cleared of terrorism, but only after having his house searched.

On August 16, a woman suffered a panic attack and became violent on a flight from London to Washington, so the plane was escorted to Boston's Logan Airport by fighter jets. "The woman was carrying hand cream and matches but was not a terrorist threat," said the TSA spokesman after the incident.

And on August 18, a plane flying from London to Egypt made an emergency landing in Italy when someone found a bomb threat scrawled on an air

sickness bag. Nothing was found on the plane, and no one knows how long the note was on board.

I'd like everyone to take a deep breath and listen for a minute.

The point of terrorism is to cause terror—sometimes to further a political goal, and sometimes out of sheer hatred. The people terrorists kill are not the targets; they are collateral damage. And blowing up planes, trains, markets, or buses is not the goal; those are just tactics. The real targets of terrorism are the rest of us: the billions of us who are not killed but are terrorized because of the killing. The real point of terrorism is not the act itself, but our reaction to the act.

And we're doing exactly what the terrorists want.

We're all a little jumpy after the recent arrest of 23 terror suspects in Great Britain. The men were reportedly plotting a liquid-explosive attack on airplanes, and both the press and politicians have been trumpeting the story ever since.

In truth, it's doubtful that their plan would have succeeded; chemists have been debunking the idea since it became public. Certainly the suspects were a long way off from trying: None had bought airline tickets, and some didn't even have passports.

Regardless of the threat, from the would-be bombers' perspective, the explosives and planes were merely tactics. Their goal was to cause terror, and in that they've succeeded.

Imagine for a moment what would have happened if they had blown up ten planes. There would be canceled flights, chaos at airports, bans on carry-on luggage, world leaders talking tough new security measures, political posturing and all sorts of false alarms as jittery people panicked. To a lesser degree, that's basically what's happening right now.

Our politicians help the terrorists every time they use fear as a campaign tactic. The press helps every time it writes scare stories about the plot and the threat. And if we're terrified, and we share that fear, we help. All of these actions intensify and repeat the terrorists' actions, and increase the effects of their terror.

(I am not saying that the politicians and press are terrorists, or that they share any of the blame for terrorist attacks. I'm not that stupid. But the subject of terrorism is more complex than it appears, and understanding its various causes and effects are vital for understanding how to best deal with it.)

The implausible plots and false alarms actually hurt us in two ways. Not only do they increase the level of fear, but they also waste time and resources

that could be better spent fighting the real threats and increasing actual security. I'll bet the terrorists are laughing at us.

Another thought experiment: Imagine for a moment that the British government had arrested the 23 suspects without fanfare. Imagine that the TSA and its European counterparts didn't engage in pointless airline security measures like banning liquids. And imagine that the press didn't write about it endlessly, and that the politicians didn't use the event to remind us all how scared we should be. If we'd reacted that way, then the terrorists would have truly failed.

It's time we calm down and fight terror with anti-terror. This does not mean that we simply roll over and accept terrorism. There are things our government can and should do to fight terrorism, most of them involving intelligence and investigation—and not focusing on specific plots.

But our job is to remain steadfast in the face of terror, to refuse to be terrorized. Our job is to not panic every time two Muslims stand together checking their watches. There are approximately 1 billion Muslims in the world, a large percentage of them not Arab, and about 320 million Arabs in the Middle East, the overwhelming majority of them not terrorists. Our job is to think critically and rationally, and to ignore the cacophony of other interests trying to use terrorism to advance political careers or increase a television show's viewership.

The surest defense against terrorism is to refuse to be terrorized. Our job is to recognize that terrorism is just one of the risks we face, and not a particularly common one at that. And our job is to fight those politicians who use fear as an excuse to take away our liberties and promote security theater that wastes money and doesn't make us any safer.

Movie-Plot Threats

Originally published in Wired, *8 September 2005*

Sometimes it seems like the people in charge of homeland security spend too much time watching action movies. They defend against specific movie plots instead of against the broad threats of terrorism.

We all do it. Our imaginations run wild with detailed and specific threats. We imagine anthrax spread from crop dusters. Or a contaminated milk supply. Or terrorist scuba divers armed with almanacs. Before long, we're envisioning an entire movie plot—without Bruce Willis to save the day. And we're scared.

Psychologically, this all makes sense. Humans have good imaginations. Box cutters and shoe bombs conjure vivid mental images. "We must protect the Super Bowl" packs more emotional punch than the vague "we should defend ourselves against terrorism."

The 9/11 terrorists used small pointy things to take over airplanes, so we ban small pointy things from airplanes. Richard Reid tried to hide a bomb in his shoes, so now we all have to take off our shoes. Recently, the Department of Homeland Security said that it might relax airplane security rules. It's not that there's a lessened risk of shoes, or that small pointy things are suddenly less dangerous. It's that those movie plots no longer capture the imagination like they did in the months after 9/11, and everyone is beginning to see how silly (or pointless) they always were.

Commuter terrorism is the new movie plot. The London bombers carried bombs into the subway, so now we search people entering the subways. They used cell phones, so we're talking about ways to shut down the cell-phone network.

It's too early to tell if hurricanes are the next movie-plot threat that captures the imagination.

The problem with movie-plot security is that it only works if we guess right. If we spend billions defending our subways, and the terrorists bomb a bus, we've wasted our money. To be sure, defending the subways makes commuting safer. But focusing on subways also has the effect of shifting attacks toward less-defended targets, and the result is that we're no safer overall.

Terrorists don't care if they blow up subways, buses, stadiums, theaters, restaurants, nightclubs, schools, churches, crowded markets or busy intersections. Reasonable arguments can be made that some targets are more attractive than others: airplanes because a small bomb can result in the death of everyone aboard, monuments because of their national significance, national events because of television coverage, and transportation because most people commute daily. But the United States is a big country; we can't defend everything.

One problem is that our nation's leaders are giving us what we want. Party affiliation notwithstanding, appearing tough on terrorism is important. Voting for missile defense makes for better campaigning than increasing intelligence funding. Elected officials want to do something visible, even if it turns out to be ineffective.

The other problem is that many security decisions are made at too low a level. The decision to turn off cell phones in some tunnels was made by those

in charge of the tunnels. Even if terrorists then bomb a different tunnel elsewhere in the country, that person did his job.

And anyone in charge of security knows that he'll be judged in hindsight. If the next terrorist attack targets a chemical plant, we'll demand to know why more wasn't done to protect chemical plants. If it targets schoolchildren, we'll demand to know why that threat was ignored. We won't accept "we didn't know the target" as an answer. Defending particular targets protects reputations and careers.

We need to defend against the broad threat of terrorism, not against specific movie plots. Security is most effective when it doesn't make arbitrary assumptions about the next terrorist act. We need to spend more money on intelligence and investigation: identifying the terrorists themselves, cutting off their funding, and stopping them regardless of what their plans are. We need to spend more money on emergency response: lessening the impact of a terrorist attack, regardless of what it is. And we need to face the geopolitical consequences of our foreign policy and how it helps or hinders terrorism.

These vague things are less visible, and don't make for good political grandstanding. But they will make us safer. Throwing money at this year's movie plot threat won't.

Fixing Intelligence Failures

Originally published in Crypto-Gram, *15 June 2002*

Could the intelligence community have connected the dots? Why didn't anyone connect the dots? How can we make sure we connect the dots next time? Dot connecting is the metaphor of the moment in Washington, as the various politicians scramble to make sure that 1) their pet ideas for improving domestic security are adopted, and 2) they don't get blamed for any dot connection failures that could have prevented 9/11.

Unfortunately, it's the wrong metaphor. We all know how to connect the dots. They're right there on the page, and they're all numbered. All you have to do is move your crayon from one dot to another, and when you're done you've drawn a lion. It's so easy a three-year-old could do it; what's wrong with the FBI and the CIA?

The problem is that the dots can only be numbered after the fact. With the benefit of hindsight, it's easy to draw lines from people in flight school here,

to secret meetings in foreign countries there, over to interesting tips from foreign governments, and then to INS records. Before 9/11, it's not so easy. Rather than thinking of intelligence as a simple connect-the-dots picture, think of it as a million unnumbered pictures superimposed on top of each other. Or a random-dot stereogram. Is it a lion, a tree, a cast iron stove, or just an unintelligible mess of dots? You try and figure it out.

This isn't to say that the United States didn't have some spectacular failures in analysis leading up to 9/11. Way back in the 30 September 2001 issue of *Crypto-Gram*, I wrote: "In what I am sure is the mother of all investigations, the CIA, NSA, and FBI have uncovered all sorts of data from their files, data that clearly indicates that an attack was being planned. Maybe it even clearly indicates the nature of the attack, or the date. I'm sure lots of information is there, in files, intercepts, computer memory." I was guessing there. It seems that there was more than I thought.

Given the bits of information that have been discussed in the press, I would have liked to think that we could have prevented this one, that there was a single Middle Eastern Terrorism desk somewhere inside the intelligence community whose job it was to stay on top of all of this. It seems that we couldn't, and that there wasn't. A budget issue, most likely.

Still, I think the "whose fault is it?" witch hunt is a bit much. Not that I mind seeing George Bush on the defensive. I've gotten sick of his "we're at war, and if you criticize me you're being unpatriotic" nonsense, and I think the enormous damage John Ashcroft has done to our nation's freedoms and liberties will take a generation and another Warren Court to fix. But all this finger-pointing between the CIA and FBI is childish, and I'm embarrassed by the Democrats who are pushing through their own poorly thought out security proposals so they're not viewed in the polls as being soft on terrorism.

My preference is for less politics and more intelligent discussion. And I'd rather see the discussion center on how to improve things for next time, rather than on who gets the blame for this time. So, in the spirit of bipartisanship (there are plenty of nitwits in both parties), here are some points for discussion:

- It's not about data collection; it's about data *analysis*. Again from the 30 September 2001 issue of *Crypto-Gram*: "Demands for even more surveillance miss the point. The problem is not obtaining data, it's deciding which data is worth analyzing and then interpreting it. Everyone already leaves a wide audit trail as we go through life, and

law enforcement can already access those records with search warrants [and subpoenas]. The FBI quickly pieced together the terrorists' identities and the last few months of their lives, once they knew where to look. If they had thrown up their hands and said that they couldn't figure out who did it or how, they might have a case for needing more surveillance data. But they didn't, and they don't.

- Security decisions need to be made as close to the source as possible. This has all sorts of implications: airport X-ray machines should be right next to the departure gates, like they are in some European airports; bomb target decisions should be made by the generals on the ground in the war zone, not by some bureaucrat in Washington; and investigation approvals should be granted the FBI office that's closest to the investigation. This mode of operation has more opportunities for abuse, so oversight is vital. But it is also more robust, and the best way to make things work. (The U.S. Marine Corps understands this principle; it's the heart of their chain of command rules.)
- Data correlation needs to happen as far away from the sources as possible. Good intelligence involves finding meaning amongst enormous reams of irrelevant data, and then organizing all those disparate pieces of information into coherent predictions about what will happen next. It requires smart people who can see connections, and access to information from many different branches of government. It can't be by the various individual pieces of bureaucracy, whether it be the CIA, FBI, NSA, INS, Coast Guard, etc. The whole picture is larger than any of them, and each one only has access to a small piece.
- Intelligence and law enforcement have fundamentally different missions. The FBI's model of operation—investigation of past crimes—does not lend itself to an intelligence paradigm: prediction of future events. On the other hand, the CIA is prohibited by law from spying on citizens. Expecting the FBI to become a domestic CIA is a terrible idea; the missions are just too different and that's too much power to consolidate under one roof. Turning the CIA into a domestic intelligence agency is an equally terrible idea; the tactics that they regularly use abroad are unconstitutional here.
- Don't forget old-fashioned intelligence gathering. Enough with the Echelon-like NSA programs where everything and anything gets sucked into an enormous electronic maw, never to be looked at again. Lots of Americans managed to become part of al-Qaeda

(a 20-year-old Californian did it, for crying out loud); why weren't any of them feeding intelligence to the CIA? Get out in the field and do your jobs.

- Organizations with investigative powers require constant oversight. If we want to formalize a domestic intelligence agency, we are going to need to be very careful about how we do it. Many of the checks and balances that Ashcroft is discarding were put in place to prevent abuse. And abuse is rampant—at the federal, state, and local levels. Just because everyone is feeling good about the police today doesn't mean that things won't change in the future. They always do.
- Fundamental changes in how the United States copes with domestic terrorism requires, um, fundamental changes. Much as the Bush administration would like to ignore the constitutional issues surrounding some of their proposals, those issues are real. Much of what the Israeli government does to combat terrorism in its country, even some of what the British government does, is unconstitutional in the United States. Security is never absolute; it always involved trade-offs. If we're going to institute domestic passports, arrest people in secret and deny them any rights, place people with Arab last names under continuous harassment, or methodically track everyone's financial dealings, we're going to have to rewrite the Constitution. At the very least, we need to have a frank and candid debate about what we're getting for what we're giving up. People might want to live in a police state, but let them at least decide willingly to live in a police state. My opinion has been that it is largely unnecessary to trade civil liberties for security, and that the best security measures—reinforcing the airplane cockpit door, putting barricades and guards around important buildings, improving authentication for telephone and Internet banking—have no effect on civil liberties. Broad surveillance is a mark of bad security.

All in all, I'm not sure how the Department of Homeland Security is going to help with any of this. Taking a bunch of ineffectual little bureaucracies and lumping them together into a single galumptious bureaucracy doesn't seem like a step in the right direction. Leaving the FBI and CIA out of the mix—the largest sources of both valuable information and turf-based problems—doesn't help, either. And if the individual organizations squabble and refuse to share information, reshuffling the chain of command isn't really going to

make any difference—it'll just add needless layers of management. And don't forget the $37 billion this is all supposed to cost, assuming there aren't the usual massive cost overruns. Couldn't we better spend that money teaching Arabic to case officers, hiring investigators, and doing various things that actually will make a difference?

The problems are about politics and policy, and not about form and structure. Fix the former, and fixing the latter becomes easy. Change the latter without fixing the former, and nothing will change.

I'm not denying the need for some domestic intelligence capability. We need something to respond to future domestic threats. I'm not happy with this conclusion, but I think it may be the best of a bunch of bad choices. Given this, the thing to do is make sure we approach that choice correctly, paying attention to constitutional protections, respecting privacy and civil liberty, and minimizing the inevitable abuses of power.

Data Mining for Terrorists

Originally published in Wired, *9 March 2006*

In the post-9/11 world, there's much focus on connecting the dots. Many believe that data mining is the crystal ball that will enable us to uncover future terrorist plots. But even in the most wildly optimistic projections, data mining isn't tenable for that purpose. We're not trading privacy for security; we're giving up privacy and getting no security in return.

Most people first learned about data mining in November 2002, when news broke about a massive government data mining program called Total Information Awareness. The basic idea was as audacious as it was repellent: Suck up as much data as possible about everyone, sift through it with massive computers, and investigate patterns that might indicate terrorist plots. Americans across the political spectrum denounced the program, and in September 2003, Congress eliminated its funding and closed its offices.

But TIA didn't die. According to *The National Journal*, it just changed its name and moved inside the Defense Department.

This shouldn't be a surprise. In May 2004, the General Accounting Office published a report that listed 122 different federal government data mining programs that used people's personal information. This list didn't include classified programs, like the NSA's eavesdropping effort, or state-run programs like MATRIX.

The promise of data mining is compelling, and convinces many. But it's wrong. We're not going to find terrorist plots through systems like this, and we're going to waste valuable resources chasing down false alarms. To understand why, we have to look at the economics of the system.

Security is always a trade-off, and for a system to be worthwhile, the advantages have to be greater than the disadvantages. A national security data mining program is going to find some percentage of real attacks, and some percentage of false alarms. If the benefits of finding and stopping those attacks outweigh the cost—in money, liberties, etc.—then the system is a good one. If not, then you'd be better off spending that cost elsewhere.

Data mining works best when there's a well-defined profile you're searching for, a reasonable number of attacks per year, and a low cost of false alarms. Credit card fraud is one of data mining's success stories: all credit card companies data mine their transaction databases, looking for spending patterns that indicate a stolen card. Many credit card thieves share a pattern—purchase expensive luxury goods, purchase things that can be easily fenced, etc.—and data mining systems can minimize the losses in many cases by shutting down the card. In addition, the cost of false alarms is only a phone call to the cardholder asking him to verify a couple of purchases. The cardholders don't even resent these phone calls—as long as they're infrequent—so the cost is just a few minutes of operator time.

Terrorist plots are different. There is no well-defined profile, and attacks are very rare. Taken together, these facts mean that data mining systems won't uncover any terrorist plots until they are very accurate, and that even very accurate systems will be so flooded with false alarms that they will be useless.

All data mining systems fail in two different ways: false positives and false negatives. A false positive is when the system identifies a terrorist plot that really isn't one. A false negative is when the system misses an actual terrorist plot. Depending on how you "tune" your detection algorithms, you can err on one side or the other: you can increase the number of false positives to ensure that you are less likely to miss an actual terrorist plot, or you can reduce the number of false positives at the expense of missing terrorist plots.

To reduce both those numbers, you need a well-defined profile. And that's a problem when it comes to terrorism. In hindsight, it was really easy to connect the 9/11 dots and point to the warning signs, but it's much harder before the fact. Certainly, there are common warning signs that many terrorist plots share, but each is unique, as well. The better you can define what you're looking

for, the better your results will be. Data mining for terrorist plots is going to be sloppy, and it's going to be hard to find anything useful.

Data mining is like searching for a needle in a haystack. There are 900 million credit cards in circulation in the United States. According to the FTC September 2003 Identity Theft Survey Report, about 1% (10 million) cards are stolen and fraudulently used each year. Terrorism is different. There are trillions of connections between people and events—things that the data mining system will have to "look at"—and very few plots. This rarity makes even accurate identification systems useless.

Let's look at some numbers. We'll be optimistic. We'll assume the system has a 1 in 100 false positive rate (99% accurate), and a 1 in 1,000 false negative rate (99.9% accurate).

Assume one trillion possible indicators to sift through: that's about ten events—e-mails, phone calls, purchases, web surfings, whatever—per person in the U.S. per day. Also assume that 10 of them are actually terrorists plotting.

This unrealistically accurate system will generate one billion false alarms for every real terrorist plot it uncovers. Every day of every year, the police will have to investigate 27 million potential plots in order to find the one real terrorist plot per month. Raise that false-positive accuracy to an absurd 99.9999% and you're still chasing 2,750 false alarms per day—but that will inevitably raise your false negatives, and you're going to miss some of those ten real plots.

This isn't anything new. In statistics, it's called the "base rate fallacy," and it applies in other domains as well. For example, even highly accurate medical tests are useless as diagnostic tools if the incidence of the disease is rare in the general population. Terrorist attacks are also rare, so any "test" is going to result in an endless stream of false alarms.

This is exactly the sort of thing we saw with the NSA's eavesdropping program: *The New York Times* reported that the computers spat out thousands of tips per month. Every one of them turned out to be a false alarm.

And the cost was enormous: not just the cost of the FBI agents running around chasing dead-end leads instead of doing things that might actually make us safer, but also the cost in civil liberties. The fundamental freedoms that make our country the envy of the world are valuable, and not something that we should throw away lightly.

Data mining can work. It helps Visa keep the costs of fraud down, just as it helps Amazon.com show me books that I might want to buy, and Google show me advertising I'm more likely to be interested in. But these are all

instances where the cost of false positives is low—a phone call from a Visa operator, or an uninteresting ad—and in systems that have value even if there is a large number of false negatives.

Finding terrorism plots is not a problem that lends itself to data mining. It's a needle-in-a-haystack problem, and throwing more hay on the pile doesn't make that problem any easier. We'd be far better off putting people in charge of investigating potential plots and letting them direct the computers, instead of putting the computers in charge and letting them decide who should be investigated.

The Architecture of Security

Originally published in Wired, *19 October 2006*

You've seen them: those large concrete blocks in front of skyscrapers, monuments, and government buildings, designed to protect against car and truck bombs. They sprang up like weeds in the months after 9/11, but the idea is much older. The prettier ones doubled as planters; the uglier ones just stood there.

Form follows function. From medieval castles to modern airports, security concerns have always influenced architecture. Castles appeared during the reign of King Stephen of England because they were the best way to defend the land and there wasn't a strong king to put any limits on castle-building. But castle design changed over the centuries in response to both innovations in warfare and politics, from motte-and-bailey to concentric design in the late medieval period to entirely decorative castles in the 19th century.

These changes were expensive. The problem is that architecture tends toward permanence, while security threats change much faster. Something that seemed a good idea when a building was designed might make little sense a century—or even a decade—later. But by then it's hard to undo those architectural decisions.

When Syracuse University built a new campus in the mid-1970s, the student protests of the late 1960s were fresh on everybody's mind. So the architects designed a college without the open greens of traditional college campuses. It's now 30 years later, but Syracuse University is stuck defending itself against an obsolete threat.

Similarly, hotel entries in Montreal were elevated above street level in the 1970s, in response to security worries about Quebecois separatists. Today the threat is gone, but those older hotels continue to be maddeningly difficult to navigate.

Also in the 1970s, the Israeli consulate in New York built a unique security system: a two-door vestibule that allowed guards to identify visitors and control building access. Now this kind of entryway is widespread, and buildings with it will remain unwelcoming long after the threat is gone.

The same thing can be seen in cyberspace as well. In his book, *Code and Other Laws of Cyberspace*, Lawrence Lessig describes how decisions about technological infrastructure—the architecture of the Internet–become embedded and then impracticable to change. Whether it's technologies to prevent file copying, limit anonymity, record our digital habits for later investigation or reduce interoperability and strengthen monopoly positions, once technologies based on these security concerns become standard it will take decades to undo them.

It's dangerously shortsighted to make architectural decisions based on the threat of the moment without regard to the long-term consequences of those decisions.

Concrete building barriers are an exception: They're removable. They started appearing in Washington, DC, in 1983, after the truck bombing of the Marines barracks in Beirut. After 9/11, they were a sort of bizarre status symbol: They proved your building was important enough to deserve protection. In New York City alone, more than 50 buildings were protected in this fashion.

Today, they're slowly coming down. Studies have found they impede traffic flow, turn into giant ashtrays, and can pose a security risk by becoming flying shrapnel if exploded.

We should be thankful they can be removed, and did not end up as permanent aspects of our cities' architecture. We won't be so lucky with some of the design decisions we're seeing about Internet architecture.

The War on the Unexpected

Originally published in Wired*, 1 November 2007*

We've opened up a new front on the war on terror. It's an attack on the unique, the unorthodox, the unexpected; it's a war on different. If you act different, you might find yourself investigated, questioned, and even arrested—even if you did nothing wrong, and had no intention of doing anything wrong. The problem is a combination of citizen informants and a CYA attitude among police that results in a knee-jerk escalation of reported threats.

This isn't the way counterterrorism is supposed to work, but it's happening everywhere. It's a result of our relentless campaign to convince ordinary citizens that they're the front line of terrorism defense. "If you see something, say something" is how the ads read in the New York City subways. "If you suspect something, report it" urges another ad campaign in Manchester, England. The Michigan State Police have a seven-minute video. Administration officials from then-attorney general John Ashcroft to DHS Secretary Michael Chertoff to President Bush have asked us all to report any suspicious activity.

The problem is that ordinary citizens don't know what a real terrorist threat looks like. They can't tell the difference between a bomb and a tape dispenser, electronic name badge, CD player, bat detector, or trash sculpture; or the difference between terrorist plotters and imams, musicians, or architects. All they know is that something makes them uneasy, usually based on fear, media hype, or just something being different.

Even worse: After someone reports a "terrorist threat," the whole system is biased towards escalation and CYA instead of a more realistic threat assessment.

Watch how it happens. Someone sees something, so he says something. The person he says it to—a policeman, a security guard, a flight attendant—now faces a choice: ignore or escalate. Even though he may believe that it's a false alarm, it's not in his best interests to dismiss the threat. If he's wrong, it'll cost him his career. But if he escalates, he'll be praised for "doing his job" and the cost will be borne by others. So he escalates. And the person he escalates to also escalates, in a series of CYA decisions. And before we're done, innocent people have been arrested, airports have been evacuated, and hundreds of police hours have been wasted.

This story has been repeated endlessly, both in the U.S. and in other countries. Someone—these are all real—notices a funny smell, or some white powder, or two people passing an envelope, or a dark-skinned man leaving boxes at the curb, or a cell phone in an airplane seat; the police cordon off the area, make arrests, and/or evacuate airplanes; and in the end the cause of the alarm is revealed as a pot of Thai chili sauce, or flour, or a utility bill, or an English professor recycling, or a cell phone in an airplane seat.

Of course, by then it's too late for the authorities to admit that they made a mistake and overreacted, that a sane voice of reason at some level should have prevailed. What follows is the parade of police and elected officials praising each other for doing a great job, and prosecuting the poor victim—the person who was different in the first place—for having the temerity to try to trick them.

For some reason, governments are encouraging this kind of behavior. It's not just the publicity campaigns asking people to come forward and snitch on

their neighbors; they're asking certain professions to pay particular attention: truckers to watch the highways, students to watch campuses, and scuba instructors to watch their students. The U.S. wanted meter readers and telephone repairmen to snoop around houses. There's even a new law protecting people who turn in their travel mates based on some undefined "objectively reasonable suspicion," whatever that is.

If you ask amateurs to act as front-line security personnel, you shouldn't be surprised when you get amateur security.

We need to do two things. The first is to stop urging people to report their fears. People have always come forward to tell the police when they see something genuinely suspicious, and should continue to do so. But encouraging people to raise an alarm every time they're spooked only squanders our security resources and makes no one safer.

We don't want people to never report anything. A store clerk's tip led to the unraveling of a plot to attack Fort Dix last May, and in March an alert Southern California woman foiled a kidnapping by calling the police about a suspicious man carting around a person-sized crate. But these incidents only reinforce the need to realistically assess, not automatically escalate, citizen tips. In criminal matters, law enforcement is experienced in separating legitimate tips from unsubstantiated fears, and allocating resources accordingly; we should expect no less from them when it comes to terrorism.

Equally important, politicians need to stop praising and promoting the officers who get it wrong. And everyone needs to stop castigating, and prosecuting, the victims just because they embarrassed the police by their innocence.

Causing a city-wide panic over blinking signs, a guy with a pellet gun, or stray backpacks, is not evidence of doing a good job: It's evidence of squandering police resources. Even worse, it causes its own form of terror, and encourages people to be even more alarmist in the future. We need to spend our resources on things that actually make us safer, not on chasing down and trumpeting every paranoid threat anyone can come up with.

Portrait of the Modern Terrorist as an Idiot

Originally published in Wired, *14 June 2007*

The recently publicized terrorist plot to blow up New York's John F. Kennedy International Airport, like so many of the terrorist plots over the past few

years, is a study in alarmism and incompetence: on the part of the terrorists, our government and the press.

Terrorism is a real threat, and one that needs to be addressed by appropriate means. But allowing ourselves to be terrorized by wannabe terrorists and unrealistic plots—and worse, allowing our essential freedoms to be lost by using them as an excuse—is wrong.

The alleged plan, to blow up JFK's fuel tanks and a small segment of the 40-mile petroleum pipeline that supplies the airport, was ridiculous. The fuel tanks are thick-walled, making them hard to damage. The airport tanks are separated from the pipelines by cutoff valves, so even if a fire broke out at the tanks, it would not back up into the pipelines. And the pipeline couldn't blow up in any case, since there's no oxygen to aid combustion. Not that the terrorists ever got to the stage—or demonstrated that they could get there—where they actually obtained explosives. Or even a current map of the airport's infrastructure.

But read what Russell Defreitas, the lead terrorist, had to say: "Anytime you hit Kennedy, it is the most hurtful thing to the United States. To hit John F. Kennedy, wow.... They love JFK—he's like the man. If you hit that, the whole country will be in mourning. It's like you can kill the man twice."

If these are the terrorists we're fighting, we've got a pretty incompetent enemy.

You couldn't tell that from the press reports, though. "The devastation that would be caused had this plot succeeded is just unthinkable," U.S. Attorney Roslynn R. Mauskopf said at a news conference, calling it "one of the most chilling plots imaginable." Sen. Arlen Specter (R-Pennsylvania) added, "It had the potential to be another 9/11."

These people are just as deluded as Defreitas.

The only voice of reason out there seemed to be New York's Mayor Michael Bloomberg, who said: "There are lots of threats to you in the world. There's the threat of a heart attack for genetic reasons. You can't sit there and worry about everything. Get a life.... You have a much greater danger of being hit by lightning than being struck by a terrorist."

And he was widely excoriated for it.

This isn't the first time a bunch of incompetent terrorists with an infeasible plot have been painted by the media as being poised to do all sorts of damage to America. In May, we learned about a six-man plan to stage an attack on Fort Dix by getting in disguised as pizza deliverymen and shooting as many soldiers and Humvees as they could, then retreating without losses to fight

again another day. Their plan, such as it was, went awry when they took a videotape of themselves at weapons practice to a store for duplication and transfer to DVD. The store clerk contacted the police, who in turn contacted the FBI. (Thank you to the video store clerk for not overreacting, and to the FBI agent for infiltrating the group.)

The "Miami 7," caught last year for plotting—among other things—to blow up the Sears Tower in Chicago, were another incompetent group: no weapons, no bombs, no expertise, no money and no operational skill. And don't forget Iyman Faris, the Ohio trucker who was convicted in 2003 for the laughable plot to take out the Brooklyn Bridge with a blowtorch. At least he eventually decided that the plan was unlikely to succeed.

I don't think these nut jobs, with their movie-plot threats, even deserve the moniker "terrorist." But in this country, while you have to be competent to pull off a terrorist attack, you don't have to be competent to cause terror. All you need to do is start plotting an attack and—regardless of whether or not you have a viable plan, weapons or even the faintest clue—the media will aid you in terrorizing the entire population.

The most ridiculous JFK Airport-related story goes to the *New York Daily News*, with its interview with a waitress who served Defreitas salmon; the front-page headline blared: "Evil Ate at Table Eight."

Following one of these abortive terror misadventures, the administration invariably jumps on the news to trumpet whatever ineffective "security" measure they're trying to push, whether it be national ID cards, wholesale National Security Agency eavesdropping (NSA), or massive data mining. Never mind that in all these cases, what caught the bad guys was old-fashioned police work—the kind of thing you'd see in decades-old spy movies.

The administration repeatedly credited the apprehension of Faris to the NSA's warrantless eavesdropping programs, even though it's just not true. The 9/11 terrorists were no different; they succeeded partly because the FBI and CIA didn't follow the leads before the attacks.

Even the London liquid bombers were caught through traditional investigation and intelligence, but this doesn't stop Secretary of Homeland Security Michael Chertoff from using them to justify access to airline passenger data.

Of course, even incompetent terrorists can cause damage. This has been repeatedly proven in Israel, and if shoe-bomber Richard Reid had been just a little less stupid and ignited his shoes in the lavatory, he might have taken out an airplane.

So these people should be locked up ... assuming they are actually guilty, that is. Despite the initial press frenzies, the actual details of the cases frequently turn out to be far less damning. Too often it's unclear whether the defendants are actually guilty, or if the police created a crime where none existed before.

The JFK Airport plotters seem to have been egged on by an informant, a twice-convicted drug dealer. An FBI informant almost certainly pushed the Fort Dix plotters to do things they wouldn't have ordinarily done. The Miami gang's Sears Tower plot was suggested by an FBI undercover agent who infiltrated the group. And in 2003, it took an elaborate sting operation involving three countries to arrest an arms dealer for selling a surface-to-air missile to an ostensible Muslim extremist. Entrapment is a very real possibility in all of these cases.

The rest of them stink of exaggeration. Jose Padilla was not actually prepared to detonate a dirty bomb in the United States, despite histrionic administration claims to the contrary. Now that the trial is proceeding, the best the government can charge him with is conspiracy to murder, kidnap, and maim, and it seems unlikely that the charges will stick. An alleged ringleader of the U.K. liquid bombers, Rashid Rauf, had charges of terrorism dropped for lack of evidence (of the 25 arrested, only 16 were charged). And now it seems the JFK mastermind was more talk than action, too.

Remember the "Lackawanna Six," those terrorists from upstate New York who pleaded guilty in 2003 to "providing support or resources to a foreign terrorist organization"? They entered their plea because they were threatened with being removed from the legal system altogether. We have no idea if they were actually guilty, or of what.

Even under the best of circumstances, these are difficult prosecutions. Arresting people before they've carried out their plans means trying to prove intent, which rapidly slips into the province of thoughtcrime. Regularly the prosecution uses obtuse religious literature in the defendants' homes to prove what they believe, and this can result in courtroom debates on Islamic theology. And then there's the issue of demonstrating a connection between a book on a shelf and an idea in the defendant's head, as if your reading of this article—or purchasing of my book—proves that you agree with everything I say. (*The Atlantic* recently published a fascinating article on this.)

I'll be the first to admit that I don't have all the facts in any of these cases. None of us does. So let's have some healthy skepticism. Skepticism when we read about these terrorist masterminds who were poised to kill thousands of

people and do incalculable damage. Skepticism when we're told that their arrest proves that we need to give away our own freedoms and liberties. And skepticism that those arrested are even guilty in the first place.

There is a real threat of terrorism. And while I'm all in favor of the terrorists' continuing incompetence, I know that some will prove more capable. We need real security that doesn't require us to guess the tactic or the target: intelligence and investigation—the very things that caught all these terrorist wannabes—and emergency response. But the "war on terror" rhetoric is more politics than rationality. We shouldn't let the politics of fear make us less safe.

Correspondent Inference Theory and Terrorism

Originally published in Wired*, 12 July 2007*

Two people are sitting in a room together: an experimenter and a subject. The experimenter gets up and closes the door, and the room becomes quieter. The subject is likely to believe that the experimenter's purpose in closing the door was to make the room quieter.

This is an example of correspondent inference theory. People tend to infer the motives—and also the disposition—of someone who performs an action based on the effects of his actions, and not on external or situational factors. If you see someone violently hitting someone else, you assume it's because he wanted to—and is a violent person—and not because he's play-acting. If you read about someone getting into a car accident, you assume it's because he's a bad driver and not because he was simply unlucky. And—more importantly for this column—if you read about a terrorist, you assume that terrorism is his ultimate goal.

It's not always this easy, of course. If someone chooses to move to Seattle instead of New York, is it because of the climate, the culture, or his career? Edward Jones and Keith Davis, who advanced this theory in the 1960s and 1970s, proposed a theory of "correspondence" to describe the extent to which this effect predominates. When an action has a high correspondence, people tend to infer the motives of the person directly from the action: e.g., hitting someone violently. When the action has a low correspondence, people tend not to make the assumption: e.g., moving to Seattle.

Like most cognitive biases, correspondent inference theory makes evolutionary sense. In a world of simple actions and base motivations, it's a good

rule of thumb that allows a creature to rapidly infer the motivations of another creature. (He's attacking me because he wants to kill me.) Even in sentient and social creatures like humans, it makes a lot of sense most of the time. If you see someone violently hitting someone else, it's reasonable to assume that he's a violent person. Cognitive biases aren't bad; they're sensible rules of thumb.

But like all cognitive biases, correspondent inference theory fails sometimes. And one place it fails pretty spectacularly is in our response to terrorism. Because terrorism often results in the horrific deaths of innocents, we mistakenly infer that the horrific deaths of innocents is the primary motivation of the terrorist, and not the means to a different end.

I found this interesting analysis in a paper by Max Abrahms in *International Security*. "Why Terrorism Does Not Work" analyzes the political motivations of 28 terrorist groups: the complete list of "foreign terrorist organizations" designated by the U.S. Department of State since 2001. He lists 42 policy objectives of those groups, and found that they only achieved them 7% of the time.

According to the data, terrorism is more likely to work if 1) the terrorists attack military targets more often than civilian ones, and 2) if they have minimalist goals like evicting a foreign power from their country or winning control of a piece of territory, rather than maximalist objectives like establishing a new political system in the country or annihilating another nation. But even so, terrorism is a pretty ineffective means of influencing policy.

There's a lot to quibble about in Abrahms' methodology, but he seems to be erring on the side of crediting terrorist groups with success. (Hezbollah's objectives of expelling both peacekeepers and Israel out of Lebanon counts as a success, but so does the "limited success" by the Tamil Tigers of establishing a Tamil state.) Still, he provides good data to support what was until recently common knowledge: Terrorism doesn't work.

This is all interesting stuff, and I recommend that you read the paper for yourself. But to me, the most insightful part is when Abrahms uses correspondent inference theory to explain why terrorist groups that primarily attack civilians do not achieve their policy goals, even if they are minimalist. Abrahms writes:

"The theory posited here is that terrorist groups that target civilians are unable to coerce policy change because terrorism has an extremely high correspondence. Countries believe that their civilian populations are attacked not because the terrorist group is protesting unfavorable external conditions such as territorial occupation or poverty. Rather, target countries infer the short-term consequences of terrorism—the deaths of innocent civilians, mass fear, loss of confidence in the government to offer protection, economic contraction, and the

inevitable erosion of civil liberties—(are) the objects of the terrorist groups. In short, target countries view the negative consequences of terrorist attacks on their societies and political systems as evidence that the terrorists want them destroyed. Target countries are understandably skeptical that making concessions will placate terrorist groups believed to be motivated by these maximalist objectives."

In other words, terrorism doesn't work, because it makes people less likely to acquiesce to the terrorists' demands, no matter how limited they might be. The reaction to terrorism has an effect completely opposite to what the terrorists want; people simply don't believe those limited demands are the actual demands.

This theory explains, with a clarity I have never seen before, why so many people make the bizarre claim that al-Qaeda terrorism—or Islamic terrorism in general—is "different": that while other terrorist groups might have policy objectives, al-Qaeda's primary motivation is to kill us all. This is something we have heard from President Bush again and again—Abrahms has a page of examples in the paper—and is a rhetorical staple in the debate.

In fact, Bin Laden's policy objectives have been surprisingly consistent. Abrahms lists four; here are six from former CIA analyst Michael Scheuer's book *Imperial Hubris*:

- End U.S. support of Israel
- Force American troops out of the Middle East, particularly Saudi Arabia
- End the U.S. occupation of Afghanistan and (subsequently) Iraq
- End U.S. support of other countries' anti-Muslim policies
- End U.S. pressure on Arab oil companies to keep prices low
- End U.S. support for "illegitimate" (i.e., moderate) Arab governments like Pakistan

Although Bin Laden has complained that Americans have completely misunderstood the reason behind the 9/11 attacks, correspondent inference theory postulates that he's not going to convince people. Terrorism, and 9/11 in particular, has such a high correspondence that people use the effects of the attacks to infer the terrorists' motives. In other words, since Bin Laden caused the death of a couple of thousand people in the 9/11 attacks, people assume that must have been his actual goal, and he's just giving lip service to what he *claims* are his goals. Even Bin Laden's actual objectives are ignored as people focus on the deaths, the destruction and the economic impact.

Perversely, Bush's misinterpretation of terrorists' motives actually helps prevent them from achieving their goals.

None of this is meant to either excuse or justify terrorism. In fact, it does the exact opposite, by demonstrating why terrorism doesn't work as a tool of persuasion and policy change. But we're more effective at fighting terrorism if we understand that it is a means to an end and not an end in itself; it requires us to understand the true motivations of the terrorists and not just their particular tactics. And the more our own cognitive biases cloud that understanding, the more we mischaracterize the threat and make bad security trade-offs.

The Risks of Cyberterrorism

Originally published in Crypto-Gram, *15 June 2003*

The threat of cyberterrorism is causing much alarm these days. We have been told to expect attacks since 9/11; that cyberterrorists would try to cripple our power system, disable air traffic control and emergency services, open dams, or disrupt banking and communications. But so far, nothing's happened. Even during the war in Iraq, which was supposed to increase the risk dramatically, nothing happened. The impending cyberwar was a big dud. Don't congratulate our vigilant security, though; the alarm was caused by a misunderstanding of both the attackers and the attacks.

These attacks are very difficult to execute. The software systems controlling our nation's infrastructure are filled with vulnerabilities, but they're generally not the kinds of vulnerabilities that cause catastrophic disruptions. The systems are designed to limit the damage that occurs from errors and accidents. They have manual overrides. These systems have been proven to work; they've experienced disruptions caused by accident and natural disaster. We've been through blackouts, telephone switch failures, and disruptions of air traffic control computers. In 1999, a software bug knocked out a nationwide paging system for a day. The results might be annoying, and engineers might spend days or weeks scrambling, but the effect on the general population has been minimal.

The worry is that a terrorist would cause a problem more serious than a natural disaster, but this kind of thing is surprisingly hard to do. Worms and viruses have caused all sorts of network disruptions, but it happened by accident. In January 2003, the SQL Slammer worm disrupted 13,000 ATMs on the

Bank of America's network. But before it happened, you couldn't have found a security expert who understood that those systems were vulnerable to that particular attack. We simply don't understand the interactions well enough to predict which kinds of attacks could cause catastrophic results, and terrorist organizations don't have that sort of knowledge either—even if they tried to hire experts.

The closest example we have of this kind of thing comes from Australia in 2000. Vitek Boden broke into the computer network of a sewage treatment plant along Australia's Sunshine Coast. Over the course of two months, he leaked hundreds of thousands of gallons of putrid sludge into nearby rivers and parks. Among the results were black creek water, dead marine life, and a stench so unbearable that residents complained. This is the only known case of someone hacking a digital control system with the intent of causing environmental harm.

Despite our predilection for calling anything "terrorism," these attacks are not. We know what terrorism is. It's someone blowing himself up in a crowded restaurant, or flying an airplane into a skyscraper. It's not infecting computers with viruses, forcing air traffic controllers to route planes manually, or shutting down a pager network for a day. That causes annoyance and irritation, not terror.

This is a difficult message for some, because these days anyone who causes widespread damage is being given the label "terrorist." But imagine for a minute the leadership of al-Qaeda sitting in a cave somewhere, plotting the next move in their jihad against the United States. One of the leaders jumps up and exclaims: "I have an idea! We'll disable their e-mail...." Conventional terrorism—driving a truckful of explosives into a nuclear power plant, for example—is still easier and much more effective.

There are lots of hackers in the world—kids, mostly—who like to play at politics and dress their own antics in the trappings of terrorism. They hack computers belonging to some other country (generally not government computers) and display a political message. We've often seen this kind of thing when two countries squabble: China vs. Taiwan, India vs. Pakistan, England vs. Ireland, U.S. vs. China (during the 2001 crisis over the U.S. spy plane that crashed in Chinese territory), the U.S. and Israel vs. various Arab countries. It's the equivalent of soccer hooligans taking out national frustrations on another country's fans at a game. It's base and despicable, and it causes real damage, but it's cyberhooliganism, not cyberterrorism.

There are several organizations that track attacks over the Internet. Over the last six months, less than 1% of all attacks originated from countries on the U.S. government's Cyber Terrorist Watch List, while 35% originated from inside the United States. Computer security is still important. People overplay the risks of cyberterrorism, but they underplay the risks of cybercrime. Fraud and espionage are serious problems. Luckily, the same countermeasures aimed at cyberterrorists will also prevent hackers and criminals. If organizations secure their computer networks for the wrong reasons, it will still be the right thing to do.