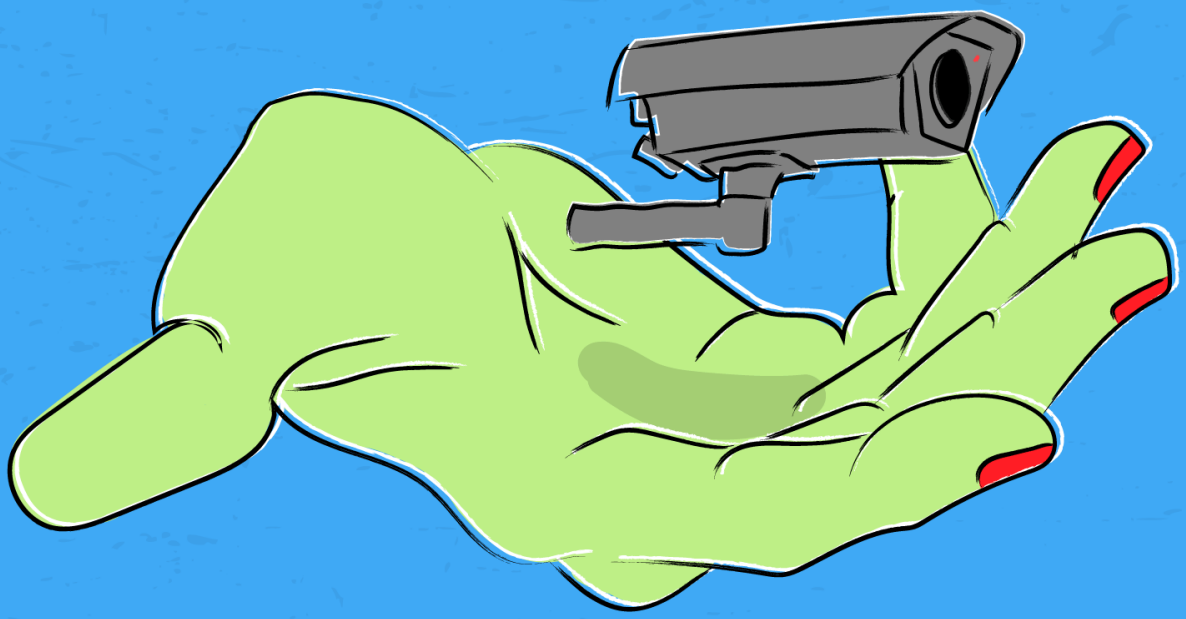

An Issue Brief on

Security Grifting

RESTORE the FOURTH

What is security grifting?

"Security grifting" is Restore the Fourth's term for a specialized subset of procurement corruption, applied to the surveillance state.

Data is a weapon in and of itself. Whoever collects it and makes it interpretable, gains power at the expense of those whose information is gathered, and their nearby communities. Government officials want to know who people are, what their thoughts are, what their movements are, because it makes them feel safer in power, not because it makes you safer.

When vendors and government agencies form a contractual relationship to gather or control data on people, without troubling about whether the data-gathering product actually does predict actual risks, that's security grifting.

Most surveillance technologies start with military uses, and it's not a secret that the Department of Defense is rife with corrupt overspending. More than two-thirds of a recent contract for "cutting edge" solutions that "directly support the warfighter" was diverted to fund efforts to lobby Congress for increased DOD funding. The tools the US government uses often "work" in the narrow tactical sense of helping make a "threat" go boom, but often fail in the larger tactical sense of correctly assessing and identifying threats, and even more often in the strategic aim of advancing US national interests.

For example, our military relies increasingly on drone warfare. Daniel Hale is in prison because he leaked documents showing that over 90% of those killed by drones were civilians. The last DOD drone strike of the Afghan war killed ten innocent members of a family, including seven children; the driver was an aid worker named Zemari Ahmadi. In the ensuing investigation, the DOD revised its description of the strike from "righteous" to a "tragic mistake," but declared that all appropriate policies had been followed. This strike wasn't exceptional. The drone was indeed working as intended. It's just that nobody within the system especially cared that the interpretive algorithm, as Mr. Ahmadi's employers put it, "could follow Zemari, an aid worker, in a commonly used car for eight hours, and not figure out who he was, and why he was at a U.S. aid organization's headquarters." Nothing in the system could accommodate to the fact that he was making stops to fill water jugs for his community, which had water shortages; nor did they pre-screen the target for the presence of children. Put simply, the system was not geared to value innocent lives.

Also in Afghanistan, the concept of "identity dominance" led to the creation of biometric databases with millions of records of Afghans working with the US government, including information on their relatives and even their favorite vegetables. This mania for data collection, with no attention paid to curation or deletion, is now proving very useful for the Taliban. Those holding the data have difficulty understanding the threat posed by someone else gaining control over it. Dutch census authorities didn't gear their systems in the 1930s against the threat that their careful gathering of data on religious affiliation would be exploited by the Nazis to facilitate genocide. The Obama administration turned a deaf ear to the many civil liberties organizations, including us, urging them to adjust executive agencies' surveillance practices to the threat of a new president weaponizing them against disfavored groups.

Let's do a thought experiment, and imagine that the FBI is considering two approaches to acquiring software to detect indicators of radicalization among social media users.

In Option A, they approach an academic specialist in the field of counterterrorism, who tells them that it is not possible to create a software product to reliably detect indicators of radicalization, because (as the academic literature amply demonstrates) there are no reliable precursors in terms of observable behavior or attitudes that can predict when someone is likely to engage in politically motivated violence.

In Option B, a vendor (let's call them "Denethor"), whose lobbyists are former senior FBI officials and members of Congress who either don't know or care about the literature, offers to develop a points-based

system for the FBI ("ThreatDTect") that will rate individuals' level of radicalization threat. It doesn't matter, for the purpose of this example, whether ThreatDTect is a facial recognition / emotion detection / micro-expressions product, a social media surveillance product, a predictive policing product based on indicia of gang association, or any of the thousand other grifts in this space.

FBI agents aren't academics. But the "security grifting" problem goes deeper than just a disregard for academic research. It's a grift because neither the vendor nor the purchaser cares primarily about whether the software operates in line with its publicly stated aims. In this example, the FBI cares primarily about gaining counter-terrorism convictions (which is subtly different from preventing politically motivated violence). A successful outcome, in terms of the FBI's incentives, is a prosecution that nets lengthy sentences, justifying both the FBI's past investments in counter-terrorism and future increases in their budget.

Considered in terms of the FBI's institutional aims, then, "ThreatDTect" will help create more successful outcomes than Option A, which could in fact undermine them. So why not show Dr. Egghead the door, kick a few tens of millions of dollars Denethor's way, and see if they deliver the goods?

On Denethor's side, it's much the same. Their incentive is to create a contract that they can renew according to the procurement schedule - ThreatDTect 2.0, 3.0, 8.0 – with just enough reasonably cheap refinements to justify the new number.

Politicians who nominally oversee the FBI also have the same incentives. Questioning a contract like this only has political costs, in terms of attack ads about being soft on terrorism. So they, too, have no incentive to care about whether it actually works.

The result is that nobody within the system cares too much whether ThreatDTect, for example, treats too many individuals as suspects without probable cause, and whether propagating ThreatDTect scores needlessly wrecks the lives of poor and vulnerable people who haven't actually done anything illegal. A high ThreatDTect score could be used in deportation proceedings, plea bargain negotiations, sentencing hearings, firearms license decisions and custody disputes, giving the government a greater whip hand in negotiations. Everybody, except the citizen (or other target of their enforcement), wins.

The same dynamic operates on the local level, with surveillance technology companies like ShotSpotter (see our forthcoming ShotSpotter brief) and the police. Police departments are often ill-equipped to evaluate what surveillance technology companies really do, and what would be appropriate limits for a given technology, in terms of data retention and access. Surveillance technology companies know this, and promise that their "black boxes" will be effective, without undertaking to comply with any formal process for demonstrating their effectiveness. They then bind police departments to NDAs, to protect their "trade secrets", and sit back to enjoy regular (and preferably increasing) contracts over time. This is part of what makes municipal and county surveillance ordinances so necessary.

Examples of security grifting

Let's go to some real-life examples from the last ten years:

In 2010, documents leaked by Anonymous and published by ProjectPM revealed the existence of a massive surveillance program targeting the Muslim world, referred to as "Romas/COIN" and then as "Odyssey." The contractors then working on the project included HB Gary, TASC, Akamai, Archimedes Global, Acclaim Technical Services, Mission Essential Personnel, Cipher, PointAbout, Google and Apple. The program was to be run out of the Department of Justice; Congress refused to investigate. It's not known whether Odyssey is still operating, what its budget is, who the contractors currently are, or whether it "works."

In 2014-2015, reports emerged of a company called DesertSnow, which trained local police in civil asset forfeiture strategies, and which provided networking software named "Black Asphalt" to

help law enforcement officers sift through publicly available data "tens of thousands of reports about American motorists, many of whom had not been charged with any crimes," including "hunches and personal data about drivers." The officer who seized the most using Black Asphalt received the honorific of "Royal Knight."

In 2016, law enforcement in San Diego pursued a contract with a "predictive policing" company called Palantir. Little is known about how the company's product works, and very little insight is given into how police officers use this type of tool and how effective it actually is. The software apparently pulls data from disparate sources, including foreclosure records and even pizza delivery, and aggregates it in one place for law enforcement.

Another example is ShadowDragon, a Wyoming company that builds software called "SocialNet" and "OIMonitor." This software is used by police departments to monitor and pull data from social media and other websites. It works by using publicly available data from social media to map out networks of individuals. The company claims the software can be used to "help predict violence and unrest" although the CEO has claimed it does not do predictive policing. It is used by CBP and ICE as well as local police. There is no good way for the public to oversee this software or whether it works, as much of ShadowDragon's operations are shrouded in secrecy.

Vendors to the incarceration industry often fit the security grifting pattern; for more details, see our forthcoming brief on "Carceral Surveillance."

Security grifting gets constantly reinforced by the merry-go-round of agency officials going to work for security grifting companies, and then returning for stints in government.

What can be done about security grifting?

Ending this scourge requires more effective laws on revolving-door appointments. A good start would be a GAO investigation of surveillance software procurement.

For any software contract with law enforcement or the intelligence community, the source code / AI + training data / benchmarks should all be available to an independent body. This body could resemble in its powers the independent Investigatory Powers Tribunal in the UK, should be attached to the Privacy and Civil Liberties Oversight Board, and should have the power to cancel contracts that cannot, at a minimum, demonstrate effectiveness in reducing threats.

The long-standing aim of a statutory charter to govern the FBI, would also discourage investigative tactics that create crime rather than detecting it.

It also requires politicians, both federally and locally, to become more systematically skeptical about the claims of vendors and the intelligence and policing professionals they work with.

Last, an aggressive public interest litigation strategy relating to the Administrative Procedures Act, could potentially result in rulings that would chill this particular kind of corruption.