THE INEFFICACY OF BIG BROTHER: ASSOCIATIONS AND THE TERROR FACTORY

The WSJ has a fascinating story, responding to (but not linking) this post, trying to address the question of whether the NSA programs we've learned about are efficient.

> But some statisticians and security experts have raised another objection: As a terror-fighting tool, it is highly inefficient and has some serious downsides.
>
> Their reasoning: Any automated approach to spotting something rare necessarily produces false positives. That means for every correctly identified target, many more alarms that go off will prove to be incorrect. So if there are vastly more innocent people than would-be terrorists whose communications are monitored, even an extremely accurate test would ensnare many non-terrorists.
>
> [snip]
>
> Even if the NSA's algorithm "is terribly clever and has a very high sensitivity and specificity, it cannot avoid having an immense false-positive rate," said Peter F. Thall, a biostatistician at the University of Texas' M.D. Anderson Cancer Center. In his arena, false positives mean patients may get tests or treatment they don't need. For the NSA, false positives could mean innocent people are monitored, detained, find themselves on no-fly lists or are otherwise inconvenienced, and that the agency spends resources inefficiently.
>
> Others, though, noted a key difference

> between terrorism and, say, a needle in
> a haystack: Terrorists tend to talk to
> each other in a way that needles don't.
> So by analyzing a network of
> communications, the NSA could be
> ferreting out clues from more than just
> the messages' particulars.

This question is, obviously, one of the reasons I posted on the 3 apparent false positives presented as implicitly terrorist associates of Najibullah Zazi in 2009. Because — assuming I'm right that they were false positives — it provides a glimpse into precisely how the government understood a lot of these terms in 2009 (I assume, though could be wrong, that their approach continues to be fine-tuned). As a reminder, here's what we know about these 3 people:

> Evidence that "individuals associated
> with Zazi purchased unusual quantities
> of hydrogen and acetone products in
> July, August, and September 2009 from
> three different beauty supply stores in
> and around Aurora;" these purchases
> include:
>
> Person one: a one-gallon container of a
> product containing 20% hydrogen peroxide
> and an 8-oz bottle of acetone
>
> Person two: an acetone product
>
> Person three: 32-oz bottles of Ion
> Sensitive Scalp Developer three
> different times

For a variety of reasons, I believe the 3 false positives consist of one person (probably person two) with a genuine relationship with Zazi who purchase relatively little acetone, and 2 people with false relationships with Zazi who bought an unusual amount of beauty supplies.

That says the FBI made two mistakes, IMO. Assuming any purchase of a common product,

acetone, was criminal on behalf of someone with a real tie to Zazi.

And assuming the relationships between the other two — the ones buying more beauty supplies — were meaningful. This could be, and I suspect it is, an assumption that anyone who belongs to the same mosque (and unlike the radical one he attended in NY, Zazi was reportedly not close to people at his mosque in CO).

Also note. This program (unlike ones I believe to exist at the National Counterterrorism Center) may not be algorithms per se at all. Rather, it could just be associations: If tie to Zazi and if beauty supply purchaser = "positive." In other words, for better and worse the FBI may not be asking the computers to "think" for it at all.

Nevertheless, the assumptions — that membership in the same mosque (or, for that matter, a single communication with a suspected terrorist) necessarily equates to a meaningful relationship — probably doom the approach in any case.

Which brings me to my other point. The WSJ suggests the costs of false positives include wasted investigative resources and unfair persecution for false positives.

But it doesn't consider the other possible uses of what may or may not be considered false positives.

First, there's the possibility an FBI investigation into a true false positive — someone totally innocent of **terrorism** — may discover some other criminal exposure, which the FBI could and has been known to use to turn the false positive into an informant.

Then there's the likelihood, especially if a potentially false positive is a young Muslim male, that the FBI will keep that person under heavy surveillance and recruitment for years and ultimately turn him into a terrorism statistic. The FBI started surveilling Mohamed Osman

Mohamud 3 years, starting before he turned 18, before they got him to attempt to bomb a public event. His parents even alerted the authorities to his increasing radicalism, but instead of intervening to reverse it, the FBI exacerbated it with several informants.

Would Mohamud have ever turned to terrorism without all that help from the FBI? Would he have developed the competence and acquired the resources to do harm? We can't actually know, and I'm actually not aware that anyone has asked this question.

What we also can't know is whether, had the FBI dedicated its efforts to something else, it could have prevented a crime developing without FBI's help.

That is, there are a whole slew of questions that have to be asked as we assess this program. Which is why we need real transparency.