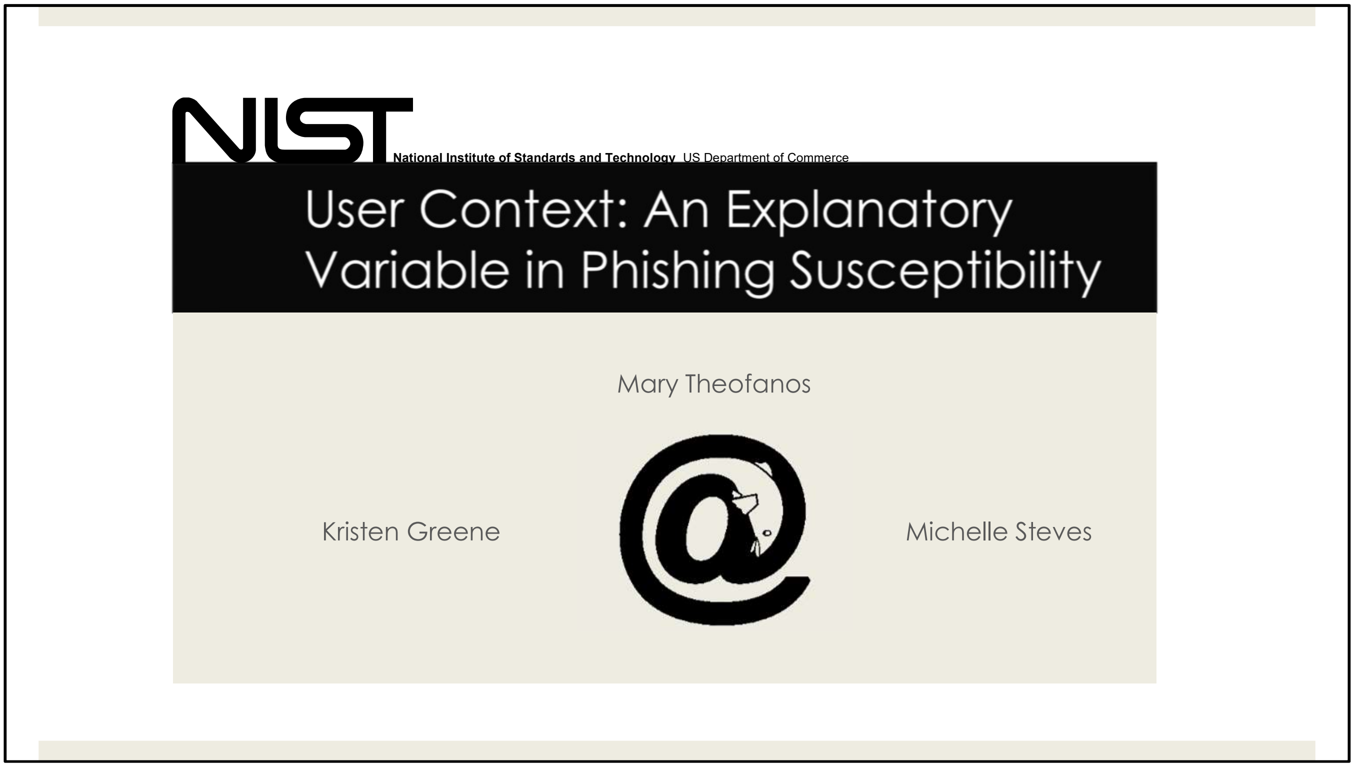

National Institute of Standards and Technology US Department of Commerce

User Context: An Explanatory Variable in Phishing Susceptibility

Mary Theofanos

Kristen Greene

Michelle Steves

Phishing has not been solved

Phishing is an ongoing issue affecting government, industry, academia, individuals, anyone who uses email is potentially affected

Linked video and press release: https://www.nist.gov/news-events/news/2018/06/youve-been-phished

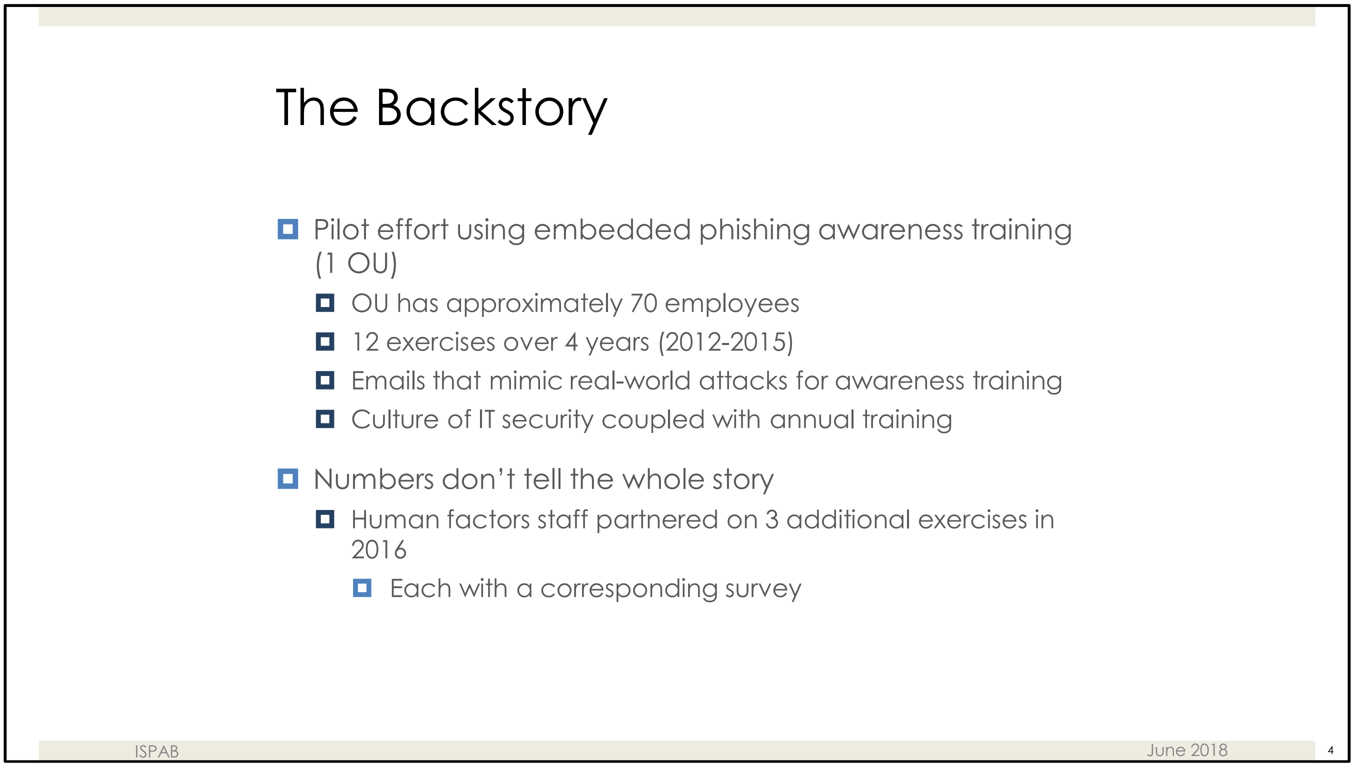

The Backstory

- Pilot effort using embedded phishing awareness training (1 OU)
  - OU has approximately 70 employees
  - 12 exercises over 4 years (2012-2015)
  - Emails that mimic real-world attacks for awareness training
  - Culture of IT security coupled with annual training
- Numbers don't tell the whole story
  - Human factors staff partnered on 3 additional exercises in 2016
    - Each with a corresponding survey

ISPAB June 2018 4

Let's talk about the research behind the video and article in more detail.

- NIST ran a pilot using embedded phishing awareness training in 1 OU for over 4 years.
- As you may know, embedded phishing awareness training consists of sending simulated phishing emails that mimic current real-world threats to staff and capturing click rates. If someone clicks, they instantly get a popup that says it was an exercise and gives info on things to look for to spot a phish.
- The OU in the pilot has approximately 70 employees.
- After 4 years of the pilot, the OISM, Office of Information Systems Management, approached our group about helping them better understand their data, and figure out what was going on with the puzzling variability in click rates they had observed.
- We partnered with OISM and the ITSO for the OU in this effort

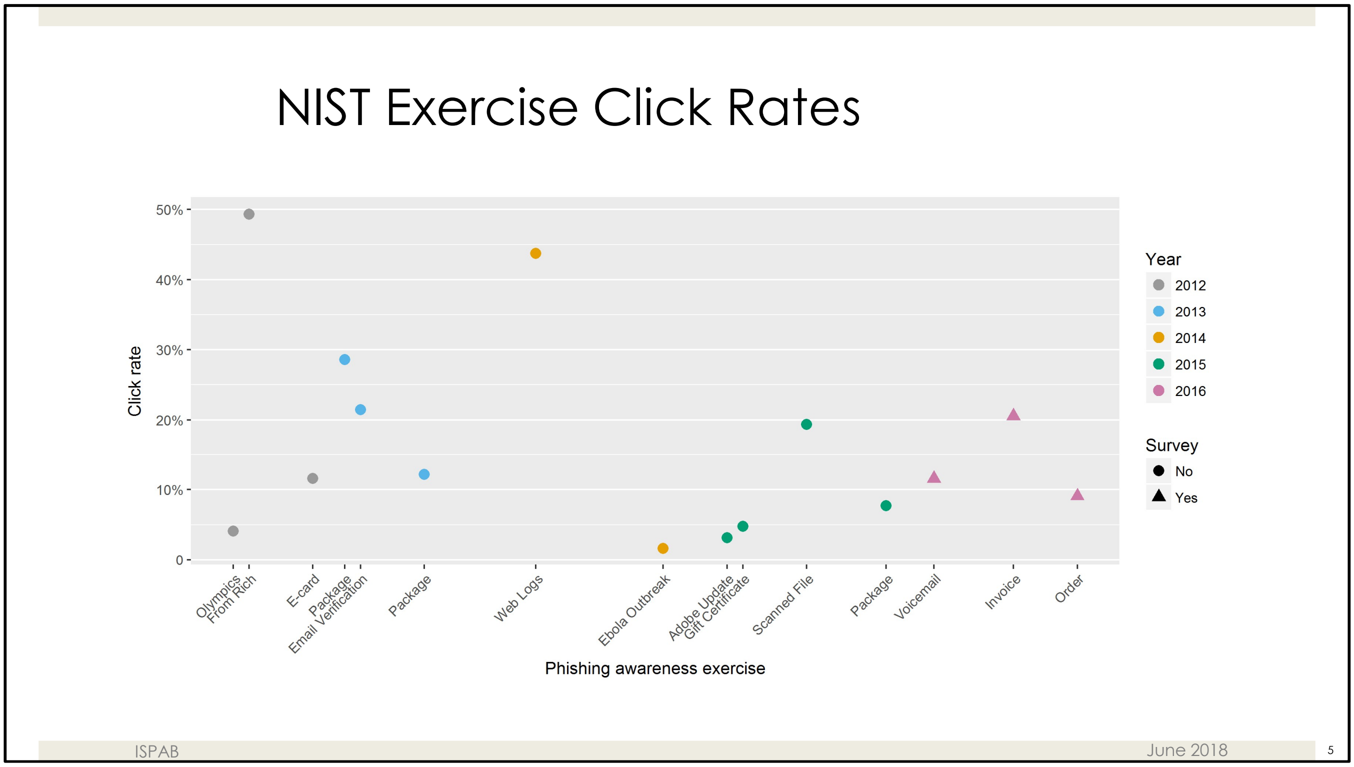

NIST Exercise Click Rates

ISPAB June 2018 5

- let's take a look at the **click rate data from the entire NIST pilot**, all 15 exercises (~1 per quarter, but with variable timing so people don't expect them)
- **X axis**, different phishing awareness scenarios that were used, ordered chorologically. **Y axis, click rates**
- **Circles represent exercises prior to our involvement;**
- We looked at the click rates and the phish scenarios; rates were variable and not going to zero (as hoped); especially note those 2 points above 40% (we'll come back to those)
- **While reviewing the data, we realized that numbers would not tell the whole story**
- **The triangles show the exercises were we participated by collecting survey data from participants.**
- They are the **last 3 scenarios shown: new voicemail, unpaid invoice, and order confirmation**

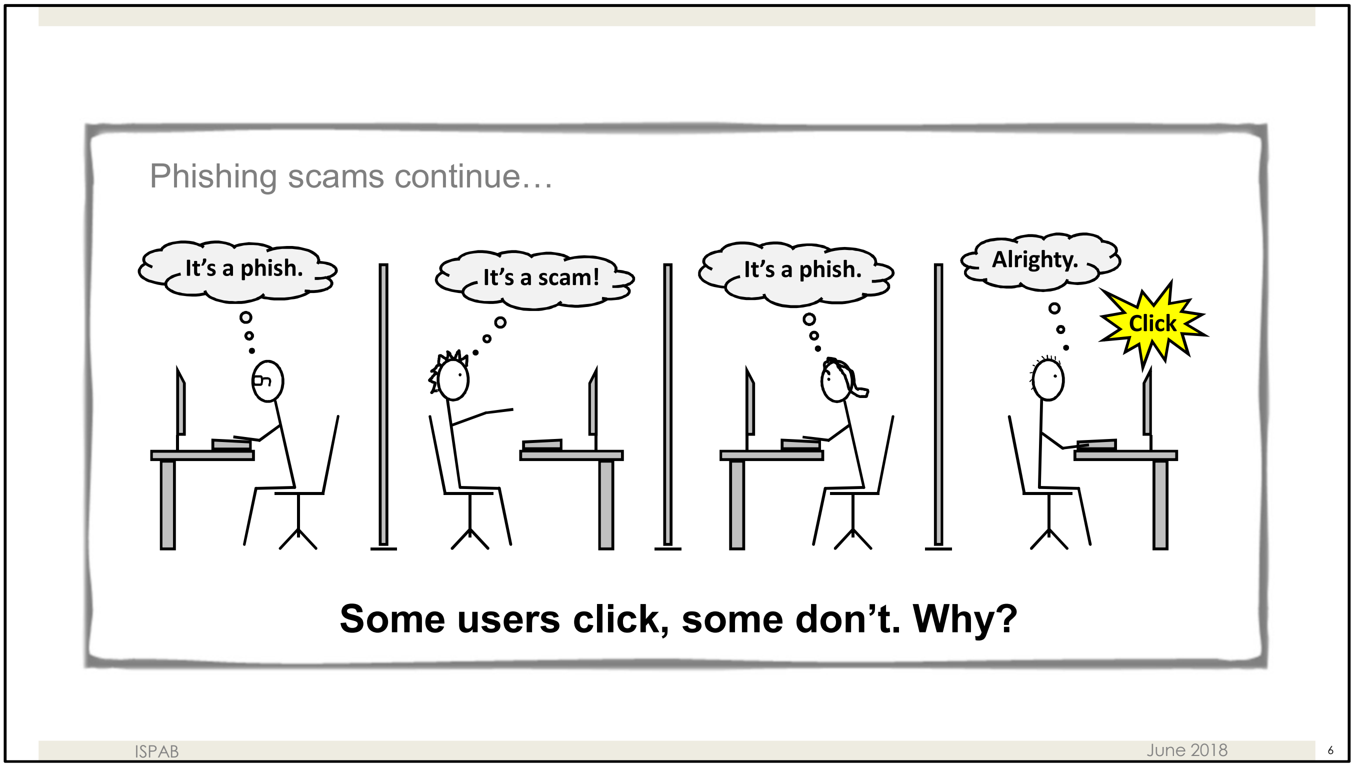

Phishing scams continue…

Some users click, some don't. Why?

ISPAB June 2018 6

Numbers are wonderful, but to dig into this, we really needed to understand from users **what's behind the numbers**—WHY do some users click when others don't?
That was our **<u>operational research question.</u>**

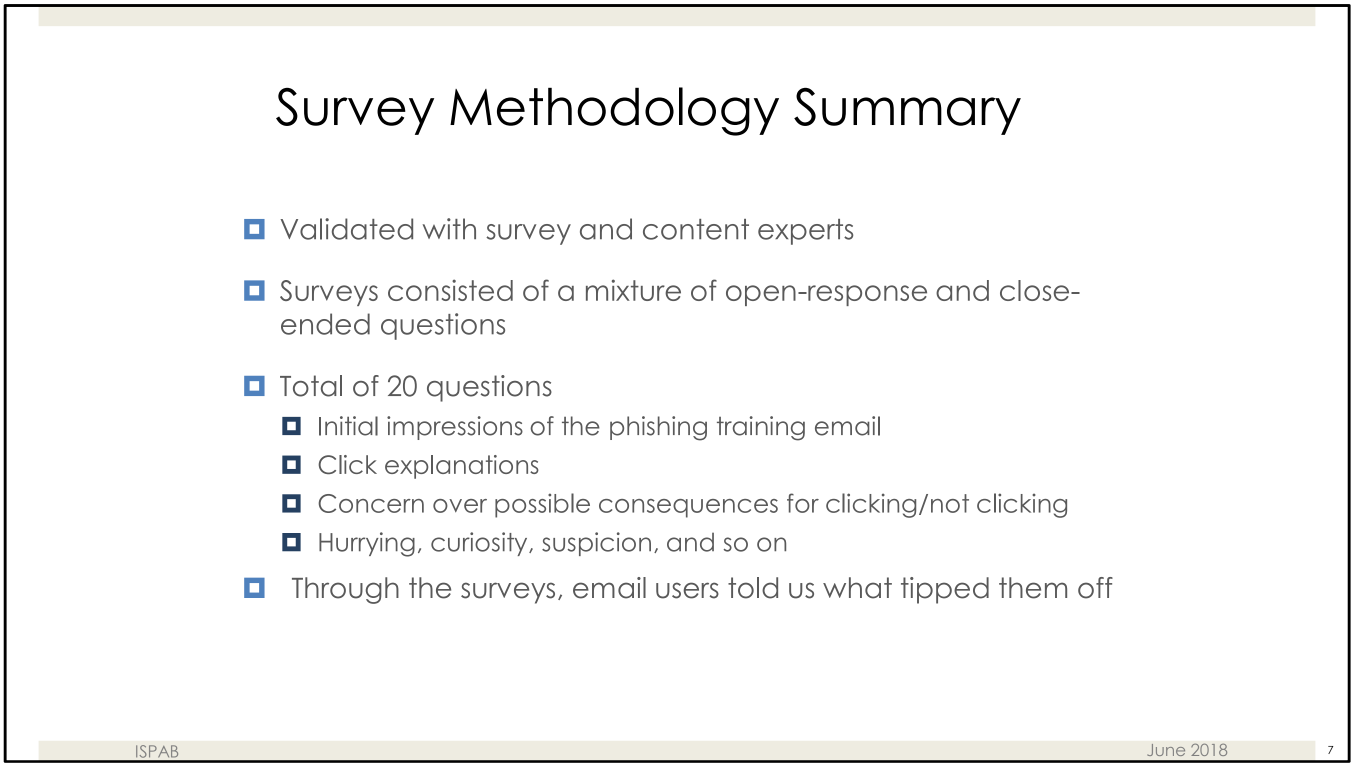

Survey Methodology Summary

- Validated with survey and content experts
- Surveys consisted of a mixture of open-response and close-ended questions
- Total of 20 questions
  - Initial impressions of the phishing training email
  - Click explanations
  - Concern over possible consequences for clicking/not clicking
  - Hurrying, curiosity, suspicion, and so on
- Through the surveys, email users told us what tipped them off

To get at this why question, chose to complement quantitative click rate data with more qualitative survey data, primarily from open-response questions….kept the survey short, followed best practices by validating with survey and content experts.
Don’t need to talk about all the survey questions, but do want to note that our survey instrument is being published and will be publicly available, so we really hope others will use it.

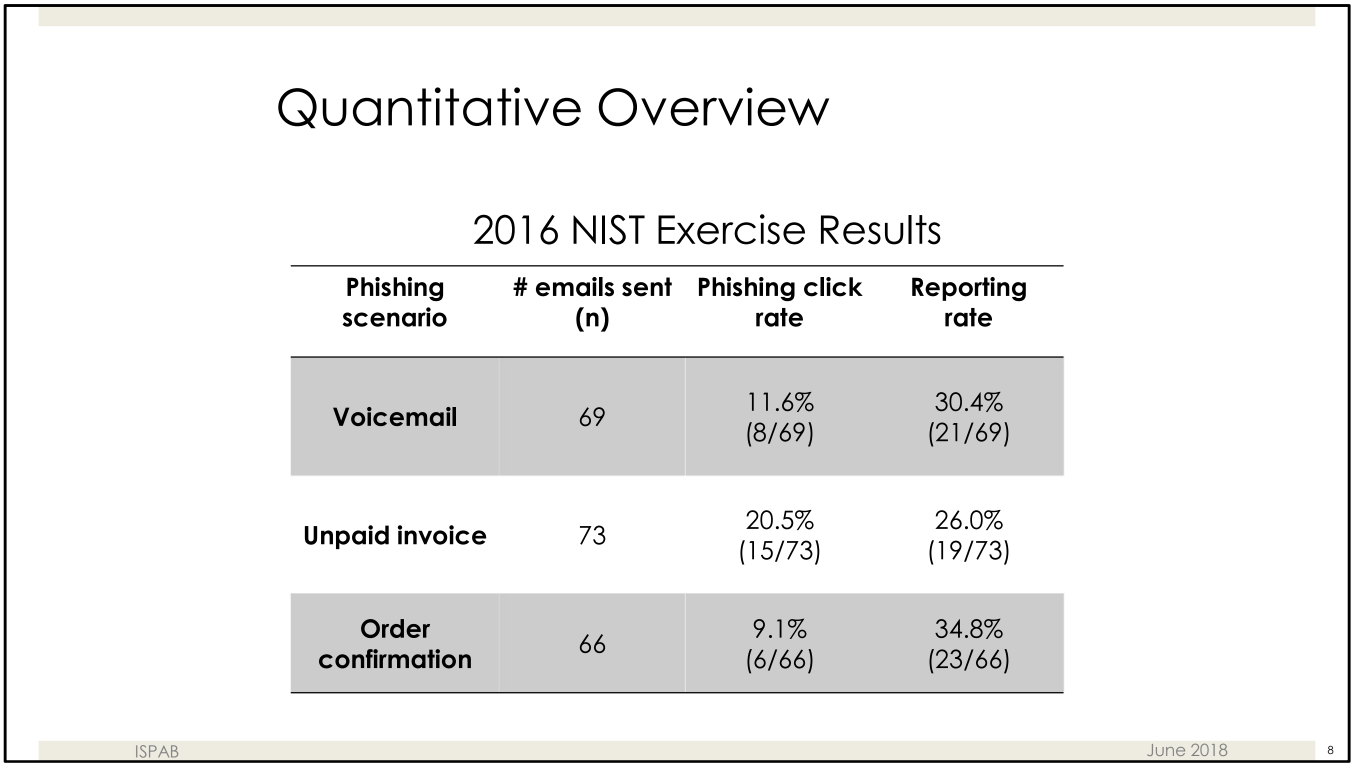

Quantitative Overview

2016 NIST Exercise Results

Phishing scenario	# emails sent (n)	Phishing click rate	Reporting rate
Voicemail	69	11.6% (8/69)	30.4% (21/69)
Unpaid invoice	73	20.5% (15/73)	26.0% (19/73)
Order confirmation	66	9.1% (6/66)	34.8% (23/66)

ISPAB June 2018 8

Before talk about the qualitative survey data, want to spend a little time on the quantitative data, both click rates and reporting rates.
Describe 3 phishing scenarios
1) New voicemail: spoofing a business process, appeared to be a system generated email, click here to listen to your new voicemail
2) Unpaid invoice: appeared to come from a fictitious Jill Preston, a fed, jill.preston@nist.gov, see the attached invoice and pay, mimicked Locky ransomware going around at the time, email said .doc was attached, but attachment was .zip
3) Order confirmation: appeared to be a system generated email confirming that an order was placed, with a link to "manage order", came out in December before the holidays

After 4 years, people are still clicking, close to 10%

- click rates for the 3 phishing awareness training exercises that we partnered with OISM on.
- **n's** ranged from **66 to 73**,
- **click rates** range from about **9% to 20%**,
- and **reporting rates** ranged from **26% to almost 35%**.

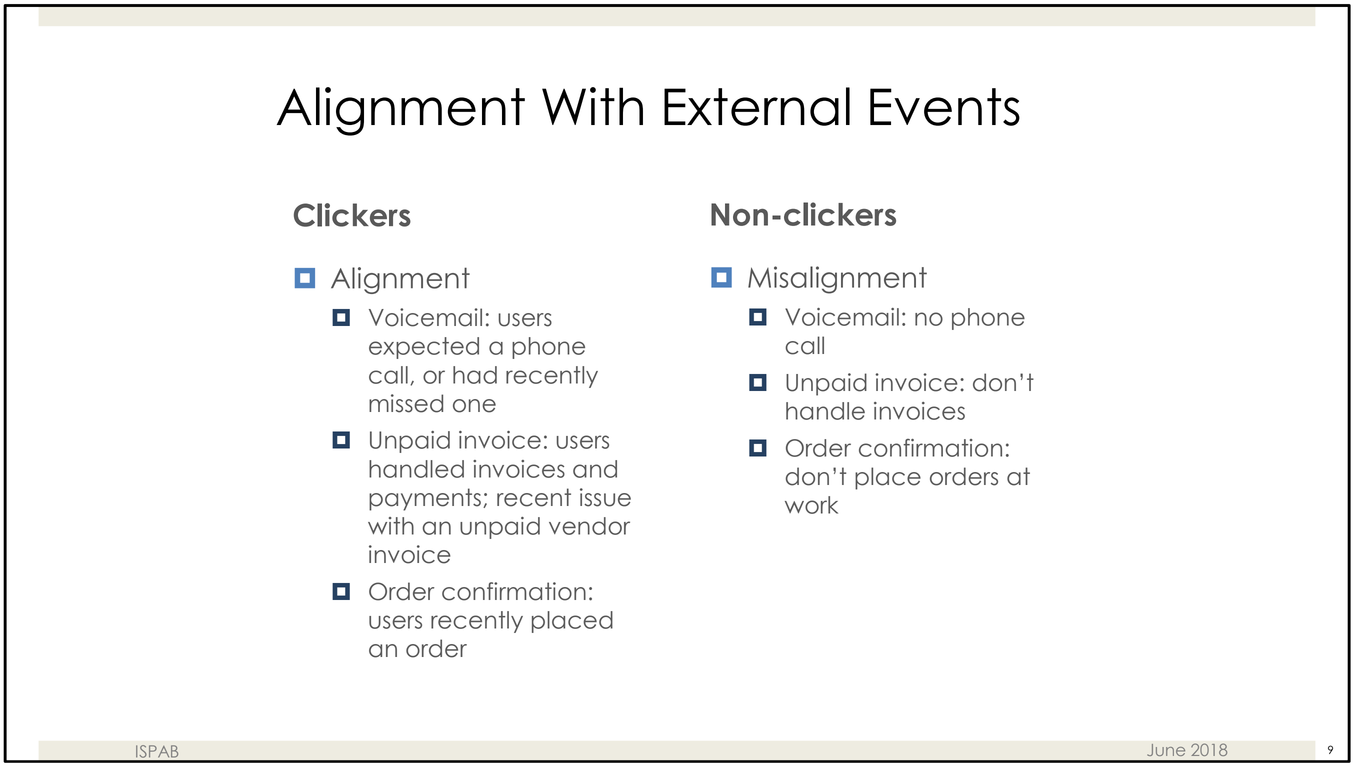

Alignment With External Events

Clickers

- Alignment
  - Voicemail: users expected a phone call, or had recently missed one
  - Unpaid invoice: users handled invoices and payments; recent issue with an unpaid vendor invoice
  - Order confirmation: users recently placed an order

Non-clickers

- Misalignment
  - Voicemail: no phone call
  - Unpaid invoice: don't handle invoices
  - Order confirmation: don't place orders at work

ISPAB June 2018 9

Click rates and reporting rates are quantitative data, transitioning to qualitative survey data now, what gives insight about what's behind the numbers
Several important themes in our qualitative data, as we step through each of them, I'll give the perspective of the clickers and contrast that with the perspective of the non-clickers, because they were very, very different. It's critical that we consider both sides of the coin; why are people clicking, and why aren't people clicking?

- Alignment with external events is one area where we say drastic differences between clickers and non-clickers
- When things are aligned, they match up. When things are misaligned, they don't.
- For clickers, the premise of the phishing email aligned with external events in their world <slide content>
- In contrast, for non-clickers, there was no such alignment—their world did NOT match the phishing scenario
- It's a user's context that determines whether there will be alignment or misalignment with the phishing scenario. By context, we mean a user's work context, their job, responsibilities, setting, etc.

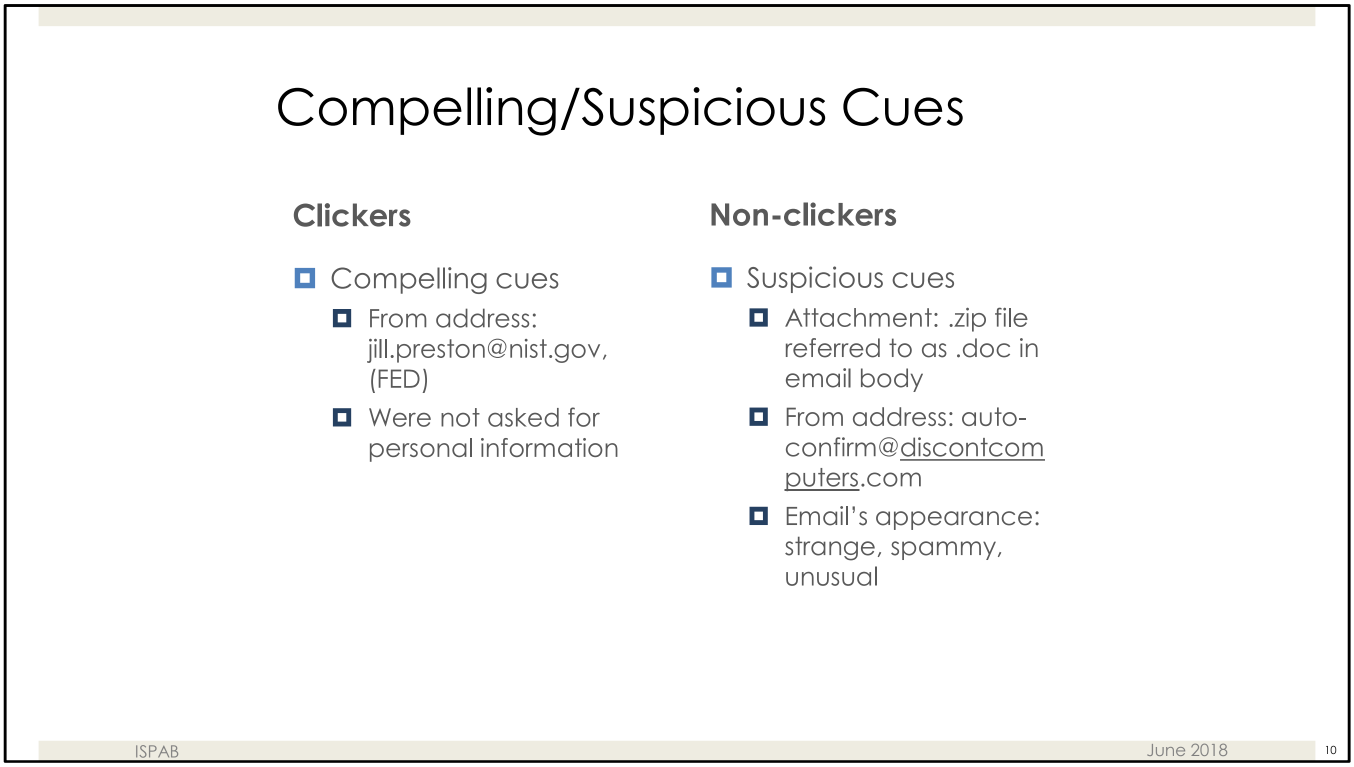

Compelling/Suspicious Cues

Clickers

- Compelling cues
  - From address: jill.preston@nist.gov, (FED)
  - Were not asked for personal information

Non-clickers

- Suspicious cues
  - Attachment: .zip file referred to as .doc in email body
  - From address: auto-confirm@discontcomputers.com
  - Email's appearance: strange, spammy, unusual

ISPAB June 2018 10

Another theme that came out in our survey data was the importance of email cues.

- **Clickers** tended to focus on compelling cues and completely ignore or discount suspicious ones.
- In the unpaid invoice phish, Jill Preston was a very compelling cue.
- Across phishing exercises, clickers found not being asked for personal info was quite compelling: they would say things like “it didn’t ask for my social or password, so I didn’t think it was a phish”. Users were not aware of the changing nature of phishing attacks, that phishing today is about more than just being asked for personal/sensitive info.
- In contrast, **non-clickers** focused on suspicious cues.
- in the unpaid invoice phish, they picked up on the attachment mismatch, where…
- In the order confirmation phish, they noticed the misspelled email address (discount)
- Across exercises, non-clickers found the email’s appearance to be suspicious. They’d say things like “my spidey sense was tingling when I saw this”

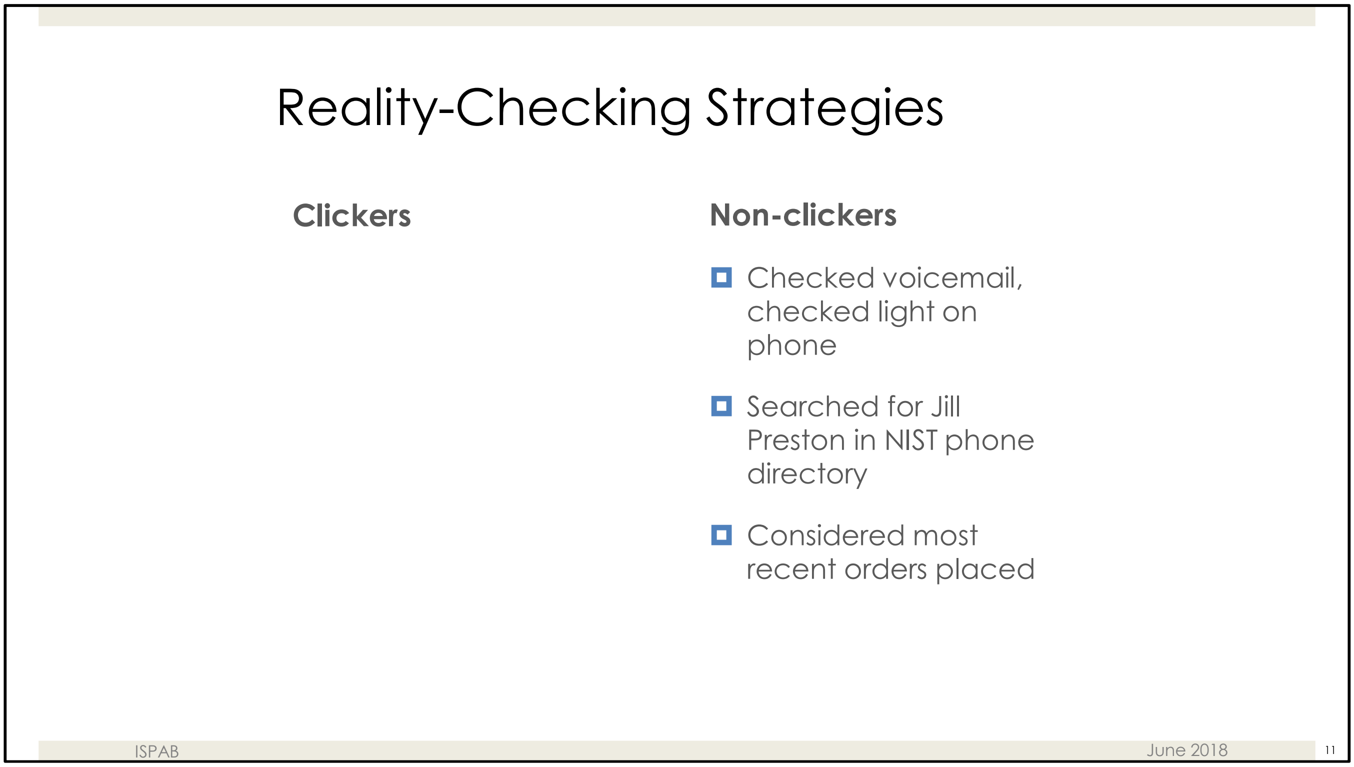

Reality-Checking Strategies

Clickers

Non-clickers

- Checked voicemail, checked light on phone
- Searched for Jill Preston in NIST phone directory
- Considered most recent orders placed

ISPAB June 2018 11

Moving on to reality-checking strategies…

- You’ll notice the clicker’s side is conspicuously blank here. That’s because reality-checking, or fact-checking strategies were something we ONLY heard from non-clickers
- they engaged in additional processing of the email in some way, engaging in deeper thought and/or actually checking some piece of information in the “real world” that could tell them whether the email was a phish.
- They really went the extra mile.
- For the voicemail phish, they…
- For the unpaid invoice phish, they…
- For the order confirmation phish, they…

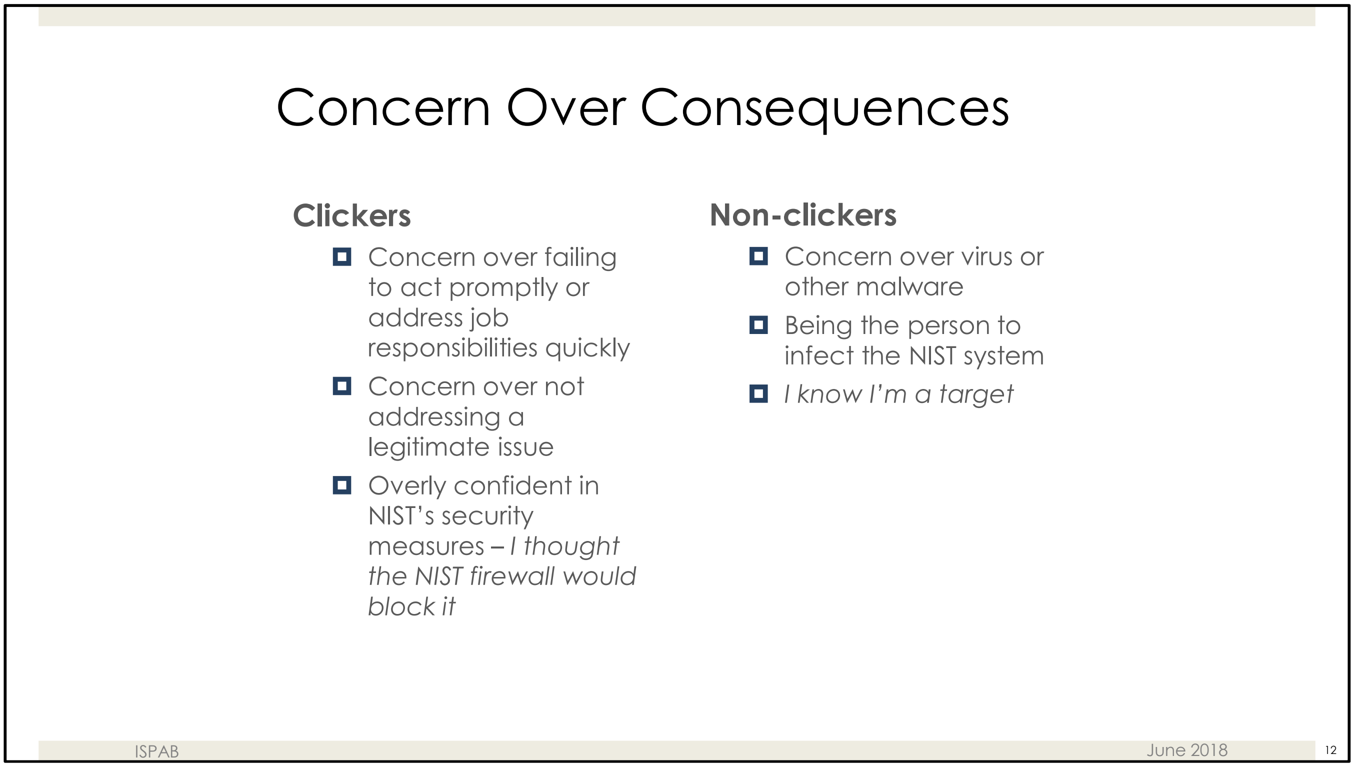

Concern Over Consequences

Clickers

- Concern over failing to act promptly or address job responsibilities quickly
- Concern over not addressing a legitimate issue
- Overly confident in NIST's security measures – *I thought the NIST firewall would block it*

Non-clickers

- Concern over virus or other malware
- Being the person to infect the NIST system
- *I know I'm a target*

Concern over consequences was another area where we saw huge differences between clickers and non-clickers
Clickers: concerned over consequences of NOT clicking. Across the different phishing exercises, they expressed concern over <slide content>

Non-clickers: concerned over consequences of clicking, they expressed concern over <slide content>

- Respondents were **much more concerned about the consequences of an unpaid invoice** than they were about a voicemail or order confirmation.
- This makes sense, since the unpaid invoice exercise happened to occur **right after a real vendor invoice** had not been paid.

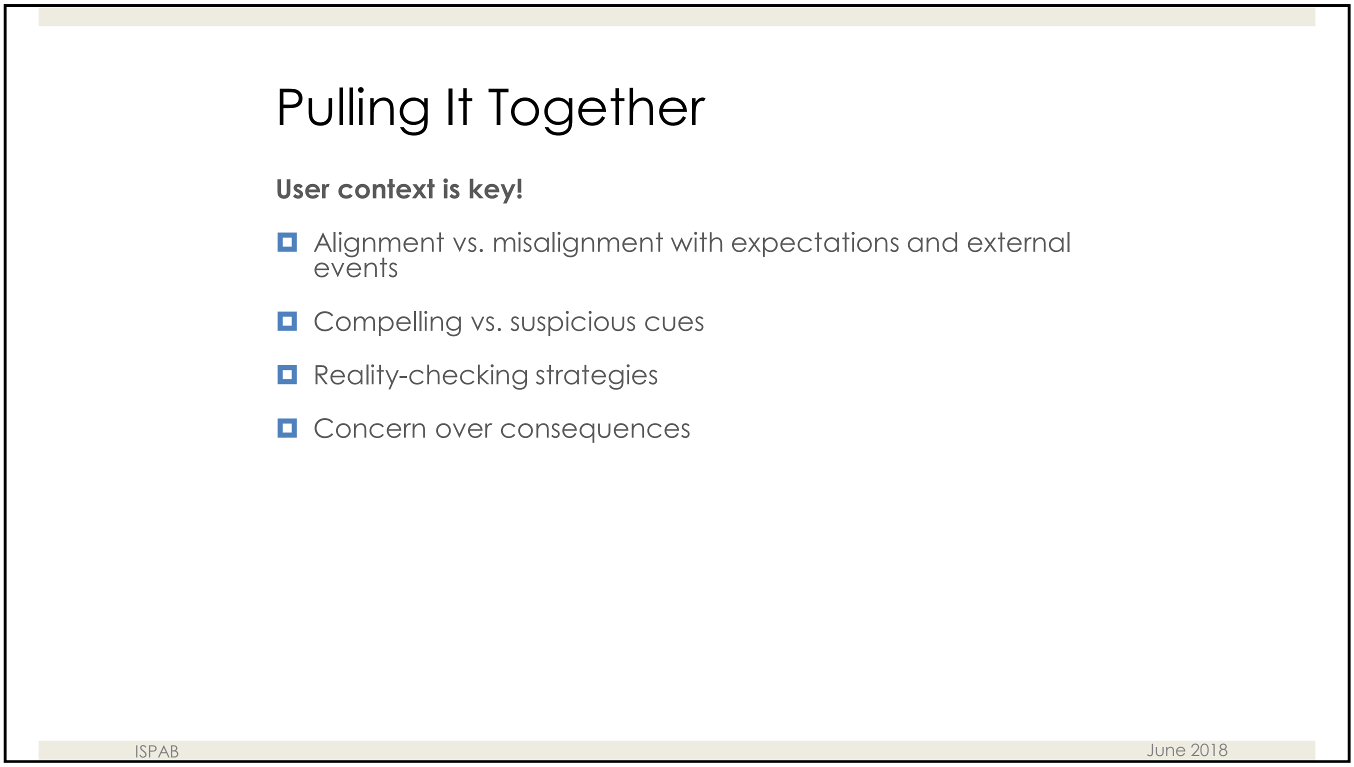

Pulling It Together

User context is key!

- Alignment vs. misalignment with expectations and external events
- Compelling vs. suspicious cues
- Reality-checking strategies
- Concern over consequences

ISPAB

June 2018

Pulling it all together, user context is key! Context is the lens through which people see and interpret the entire email

Looking across all exercises…

Understanding how a user’s context influences their clicking behavior was our “ah-ha” moment. That is the theme that ties all of our results together….it all ties back to a user’s context. In this case, what’s their job, what are their responsibilities, what have they done recently…

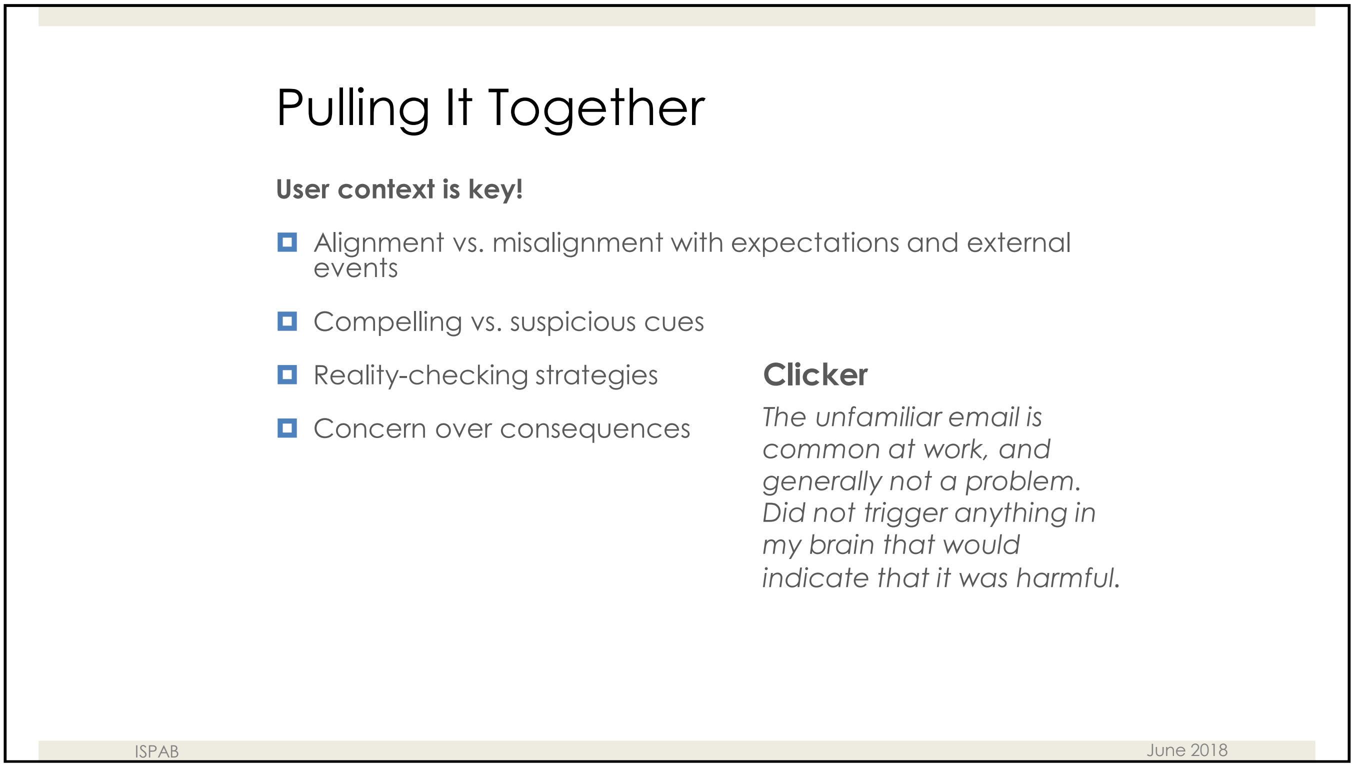

Pulling It Together

User context is key!

- Alignment vs. misalignment with expectations and external events
- Compelling vs. suspicious cues
- Reality-checking strategies
- Concern over consequences

Clicker

The unfamiliar email is common at work, and generally not a problem. Did not trigger anything in my brain that would indicate that it was harmful.

To illustrate, clickers would say things like…

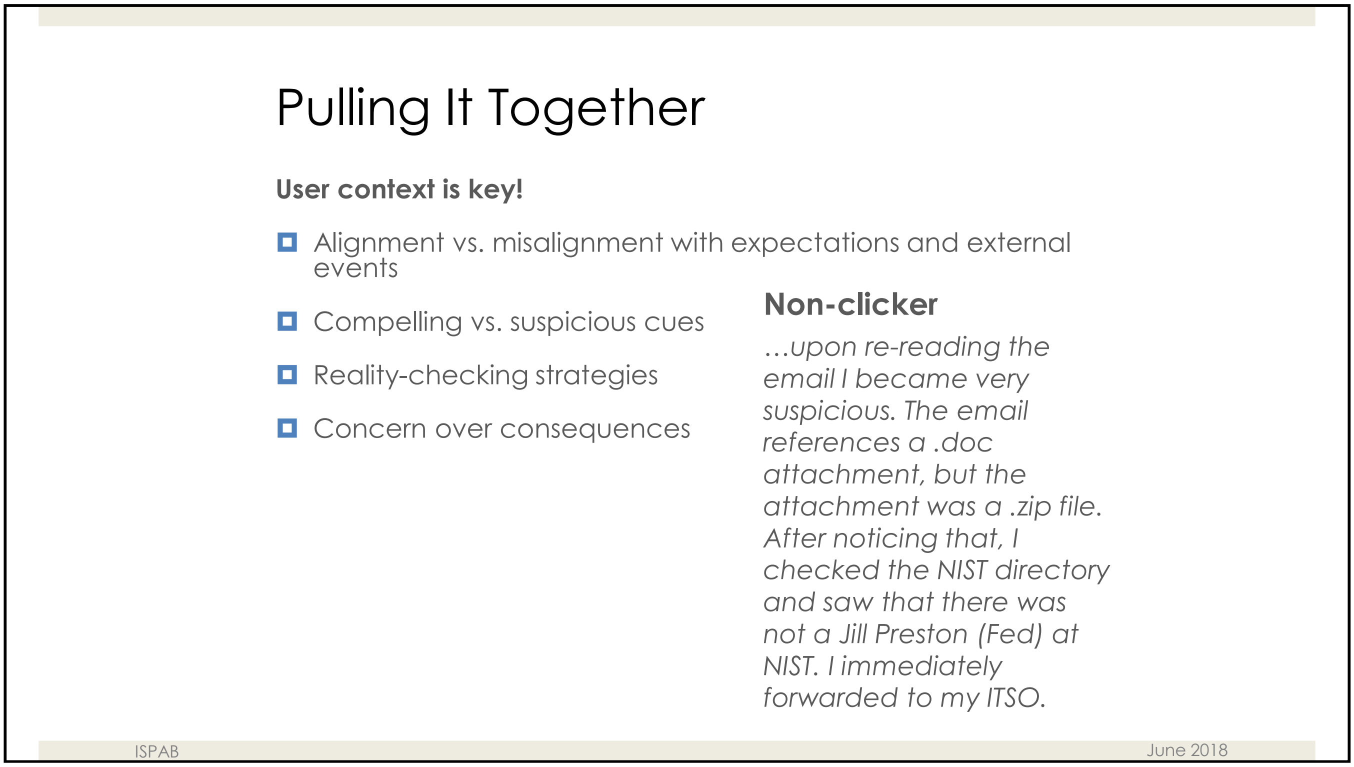

Pulling It Together

User context is key!

- Alignment vs. misalignment with expectations and external events
- Compelling vs. suspicious cues
- Reality-checking strategies
- Concern over consequences

Non-clicker

…upon re-reading the email I became very suspicious. The email references a .doc attachment, but the attachment was a .zip file. After noticing that, I checked the NIST directory and saw that there was not a Jill Preston (Fed) at NIST. I immediately forwarded to my ITSO.

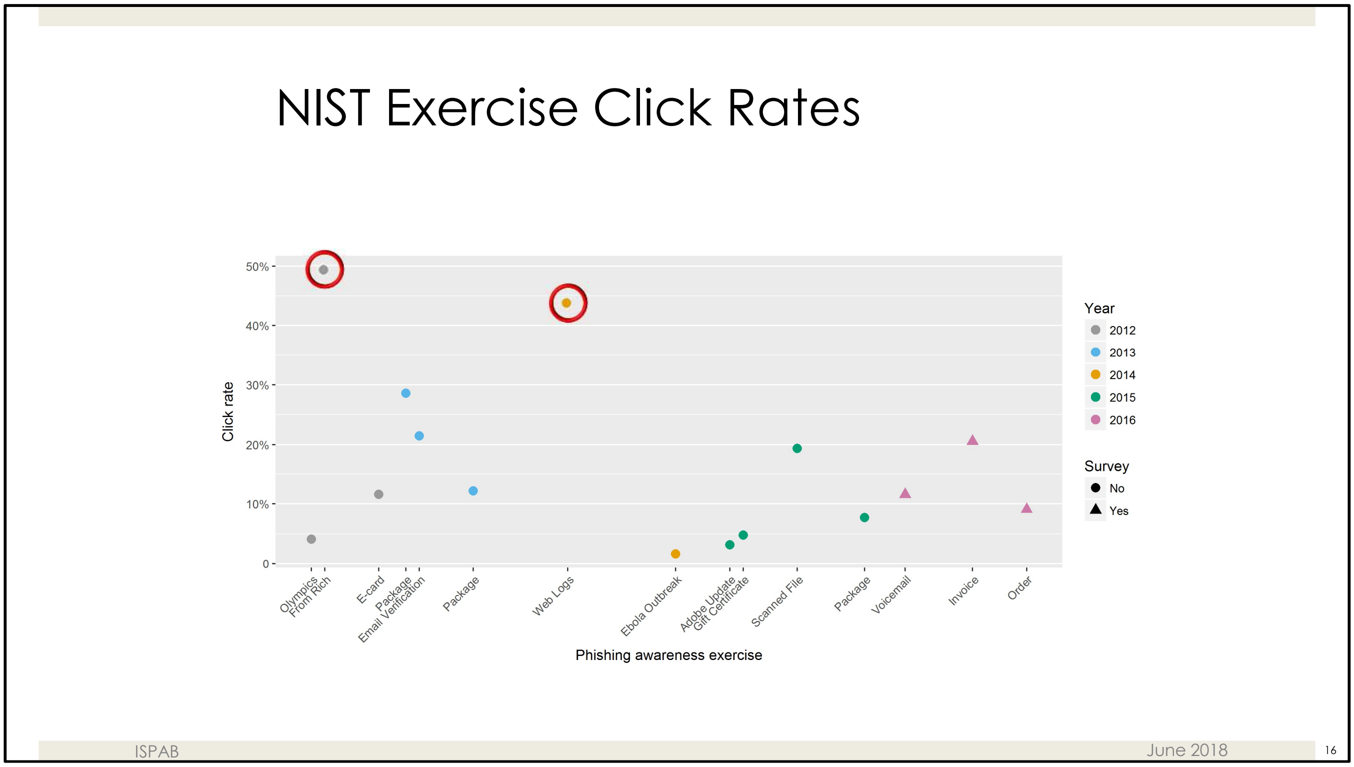

NIST Exercise Click Rates

We said before that we would come back to these points; based on our survey data, now we understand these click rates were elevated because phishing premise aligned with user context for more people in these exercises. You see the importance of context more in well-tailored phishes. These exercises were targeted at FEDs, in particular they were tailored for those who work at NIST, high click rate of ~50%.

Heard from another government agency, NRC, that their phishing exercises targeted to FEDs also have ~50% click rate—and that agency has been doing **5 exercises per quarter for almost 4 years** . Their click rates are comparable to ours, both in terms of spiking like this, where click rates jump up to 50% for a well-tailored email, and with respect to the fact that their click rates still aren't zero either

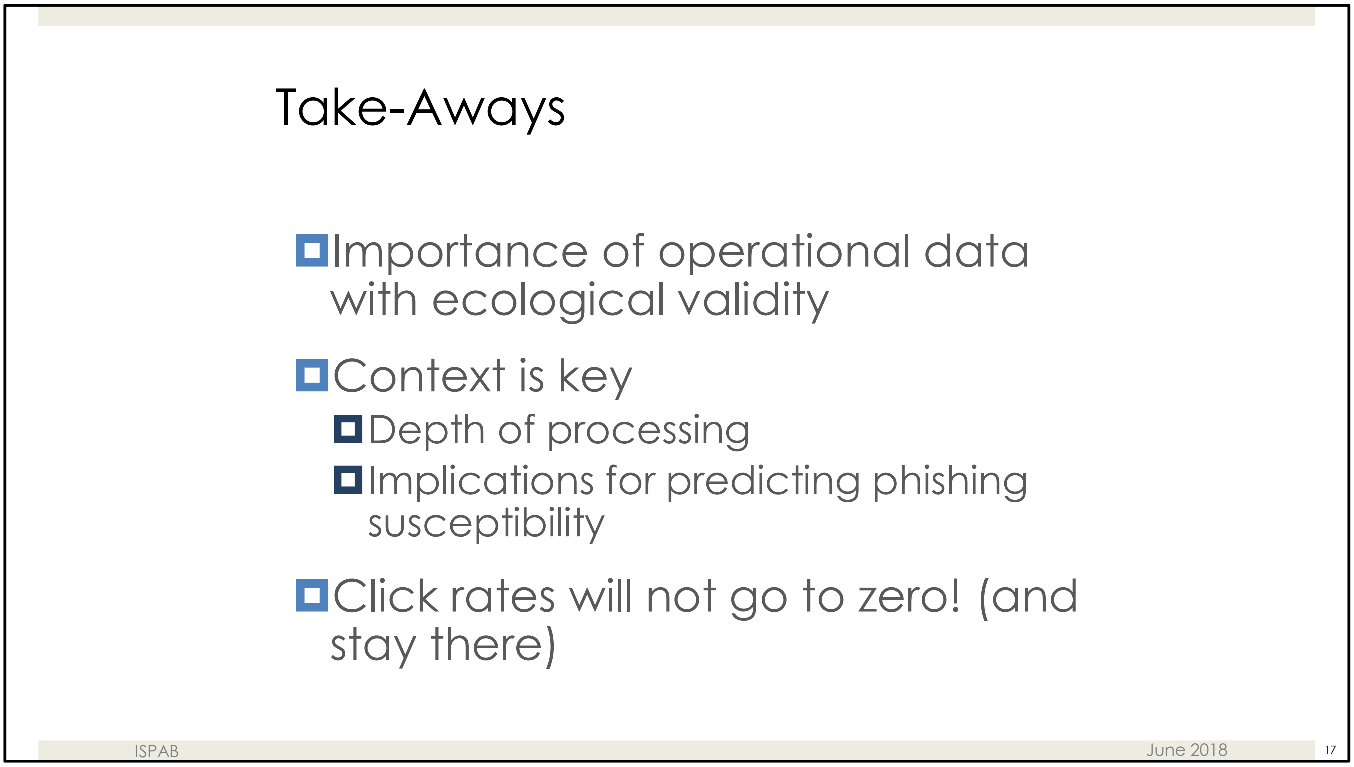

Take-Aways

- Importance of operational data with ecological validity
- Context is key
  - Depth of processing
  - Implications for predicting phishing susceptibility
- Click rates will not go to zero! (and stay there)

ISPAB June 2018 17

We can not underestimate the value of ecologically valid data. Here real federal workers are in their regular work environments, processing their normal emails with their typical work assignments and time pressures. Even though these staff members knew they were in the pilot, they did not know when the exercises would be conducted (no priming effects).

- That’s what makes our work SOOOOO different from most phishing research, which is often conducted in laboratory or university settings.
- This setting allowed the explanatory variable of user context to surface, this is certainly related to depth of processing.
- These findings have real implications for predicting phishing susceptibility – click rates will not go to zero, the context will align for someone.

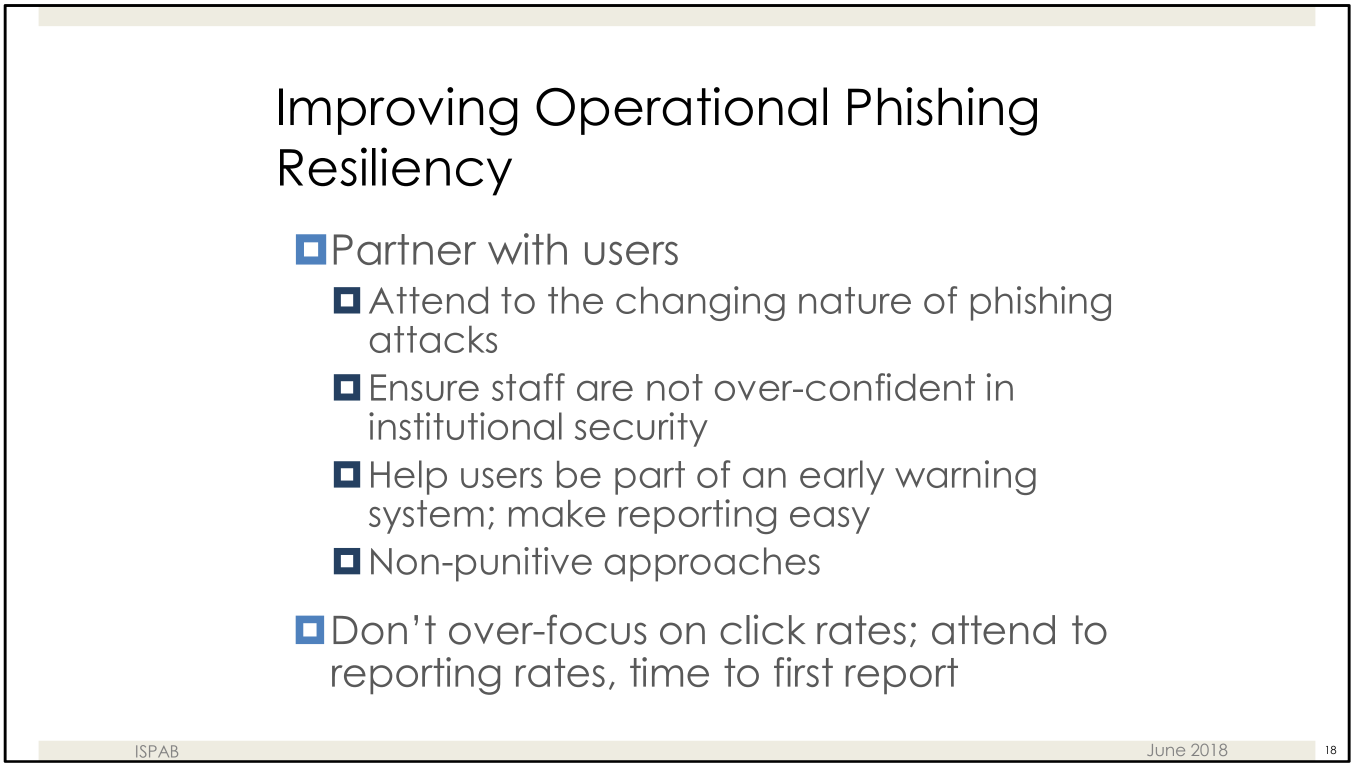

Improving Operational Phishing Resiliency

- Partner with users
  - Attend to the changing nature of phishing attacks
  - Ensure staff are not over-confident in institutional security
  - Help users be part of an early warning system; make reporting easy
  - Non-punitive approaches
- Don't over-focus on click rates; attend to reporting rates, time to first report

ISPAB June 2018 18

Technological solutions are reactive; humans can not spot every phish. How do we improve phishing resiliency?

- Engage and partner with users…

And you have to use and leverage your reporting metrics – mitigating damage and reducing clean-up is less costly in the long run
Other metrics: what types to phish and spear phish are being used against your org? Is your training reflecting those threats and emerging threats?

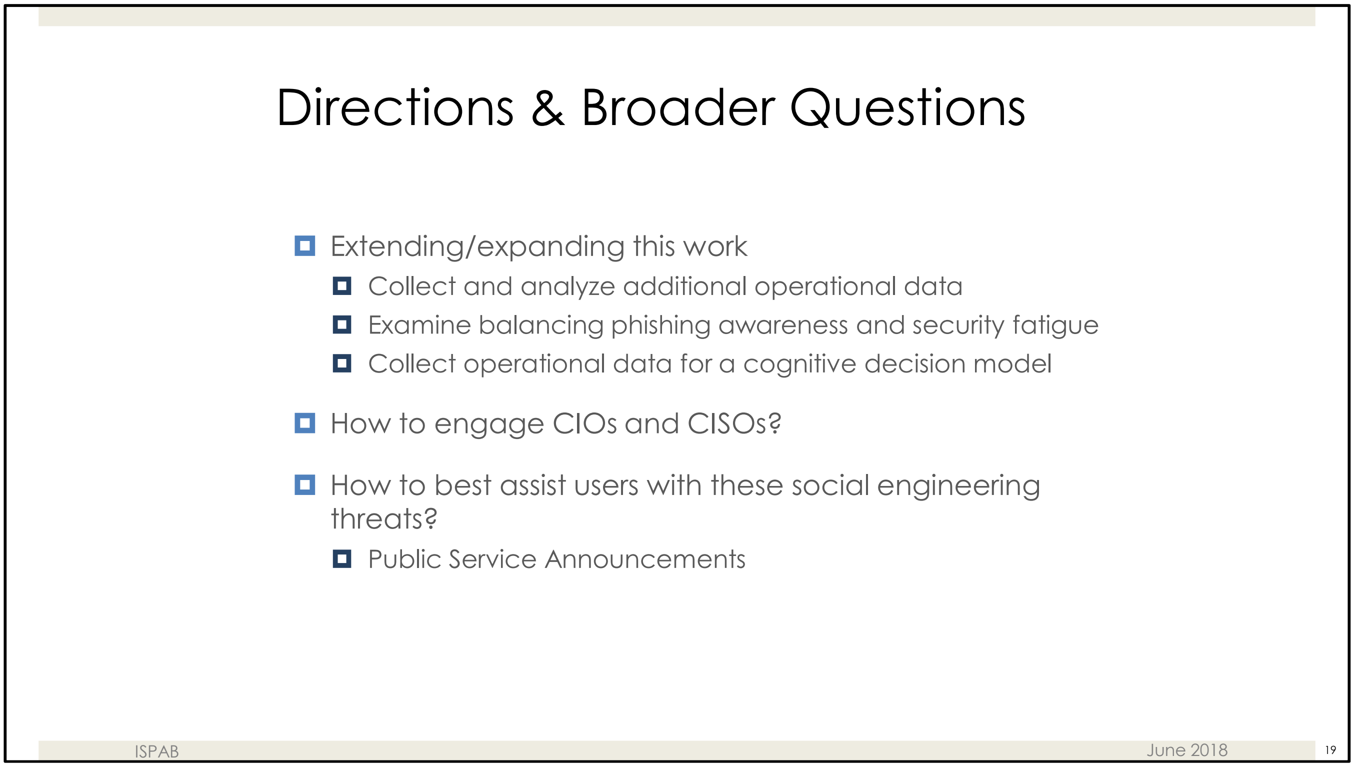

Directions & Broader Questions

- Extending/expanding this work
  - Collect and analyze additional operational data
  - Examine balancing phishing awareness and security fatigue
  - Collect operational data for a cognitive decision model
- How to engage CIOs and CISOs?
- How to best assist users with these social engineering threats?
  - Public Service Announcements

Balancing phishing awareness and security fatigue

- If we did a phishing training exercise right after IT security day, would people report more highly b/c they were more aware…(heard anecdotally from one reviewer that reporting goes way up after training then fades—what does that curve look like, is it even possible to stay aware all the time?)
  Balancing security fatigue vs. awareness
  How often should we train people
  How often should CIO put out emails on phishing threats, topic of the day

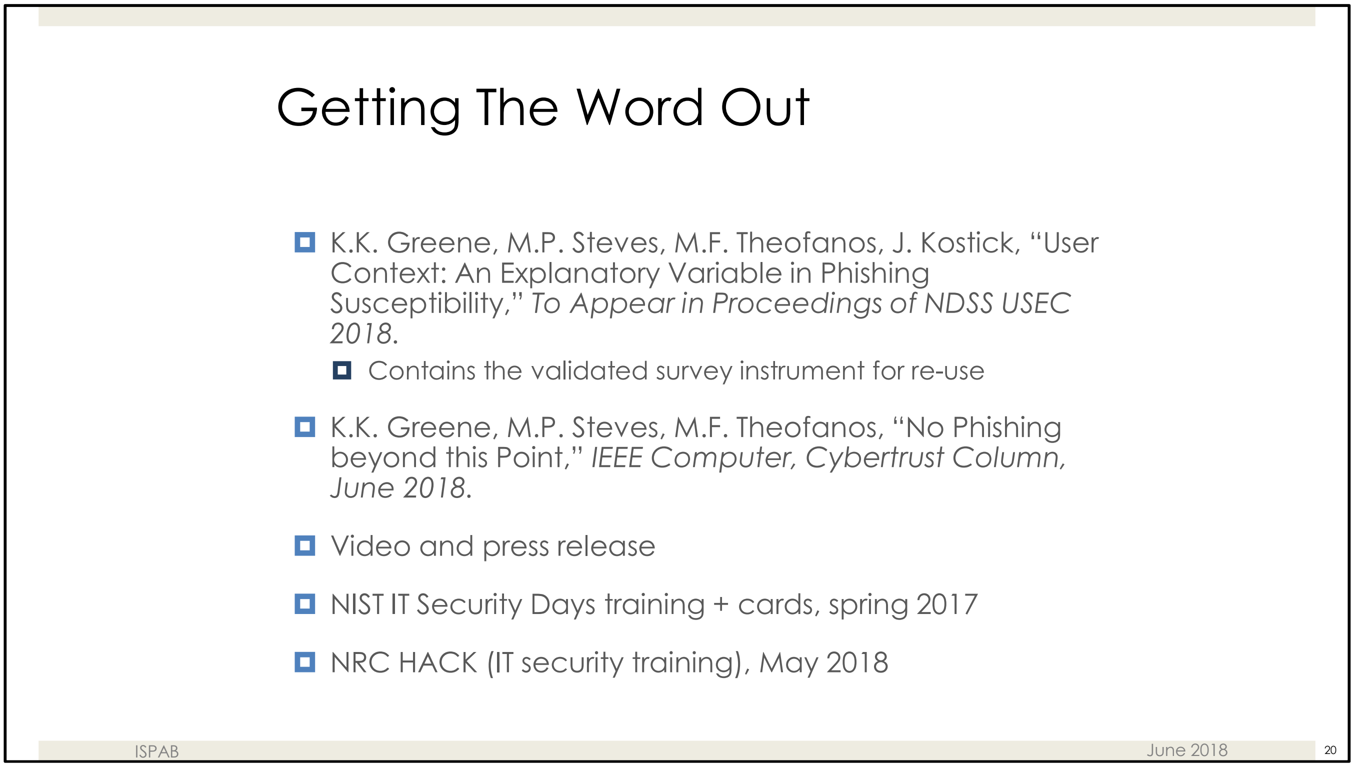

Getting The Word Out

- K.K. Greene, M.P. Steves, M.F. Theofanos, J. Kostick, "User Context: An Explanatory Variable in Phishing Susceptibility," *To Appear in Proceedings of NDSS USEC 2018.*
  - Contains the validated survey instrument for re-use
- K.K. Greene, M.P. Steves, M.F. Theofanos, "No Phishing beyond this Point," *IEEE Computer, Cybertrust Column, June 2018.*
- Video and press release
- NIST IT Security Days training + cards, spring 2017
- NRC HACK (IT security training), May 2018

Thank You

kgreene@nist.gov
msteves@nist.gov
maryt@nist.gov

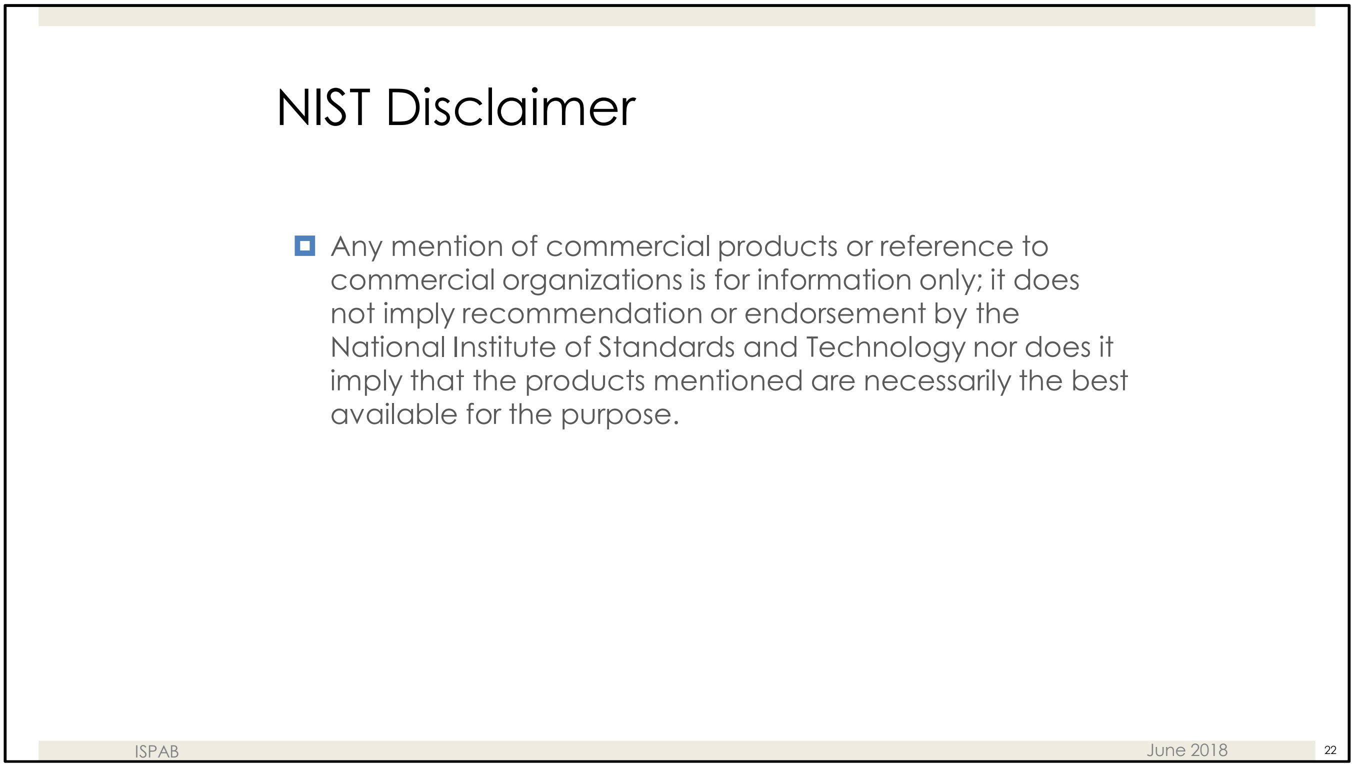

NIST Disclaimer

- Any mention of commercial products or reference to commercial organizations is for information only; it does not imply recommendation or endorsement by the National Institute of Standards and Technology nor does it imply that the products mentioned are necessarily the best available for the purpose.