Empiricism in Security

CS499/579 :: Empirical Computer Security

Zane Ma (he/him/his)
2023.09.27

IT comes for us all…

OSU's security policy requires a password change every 180 days.

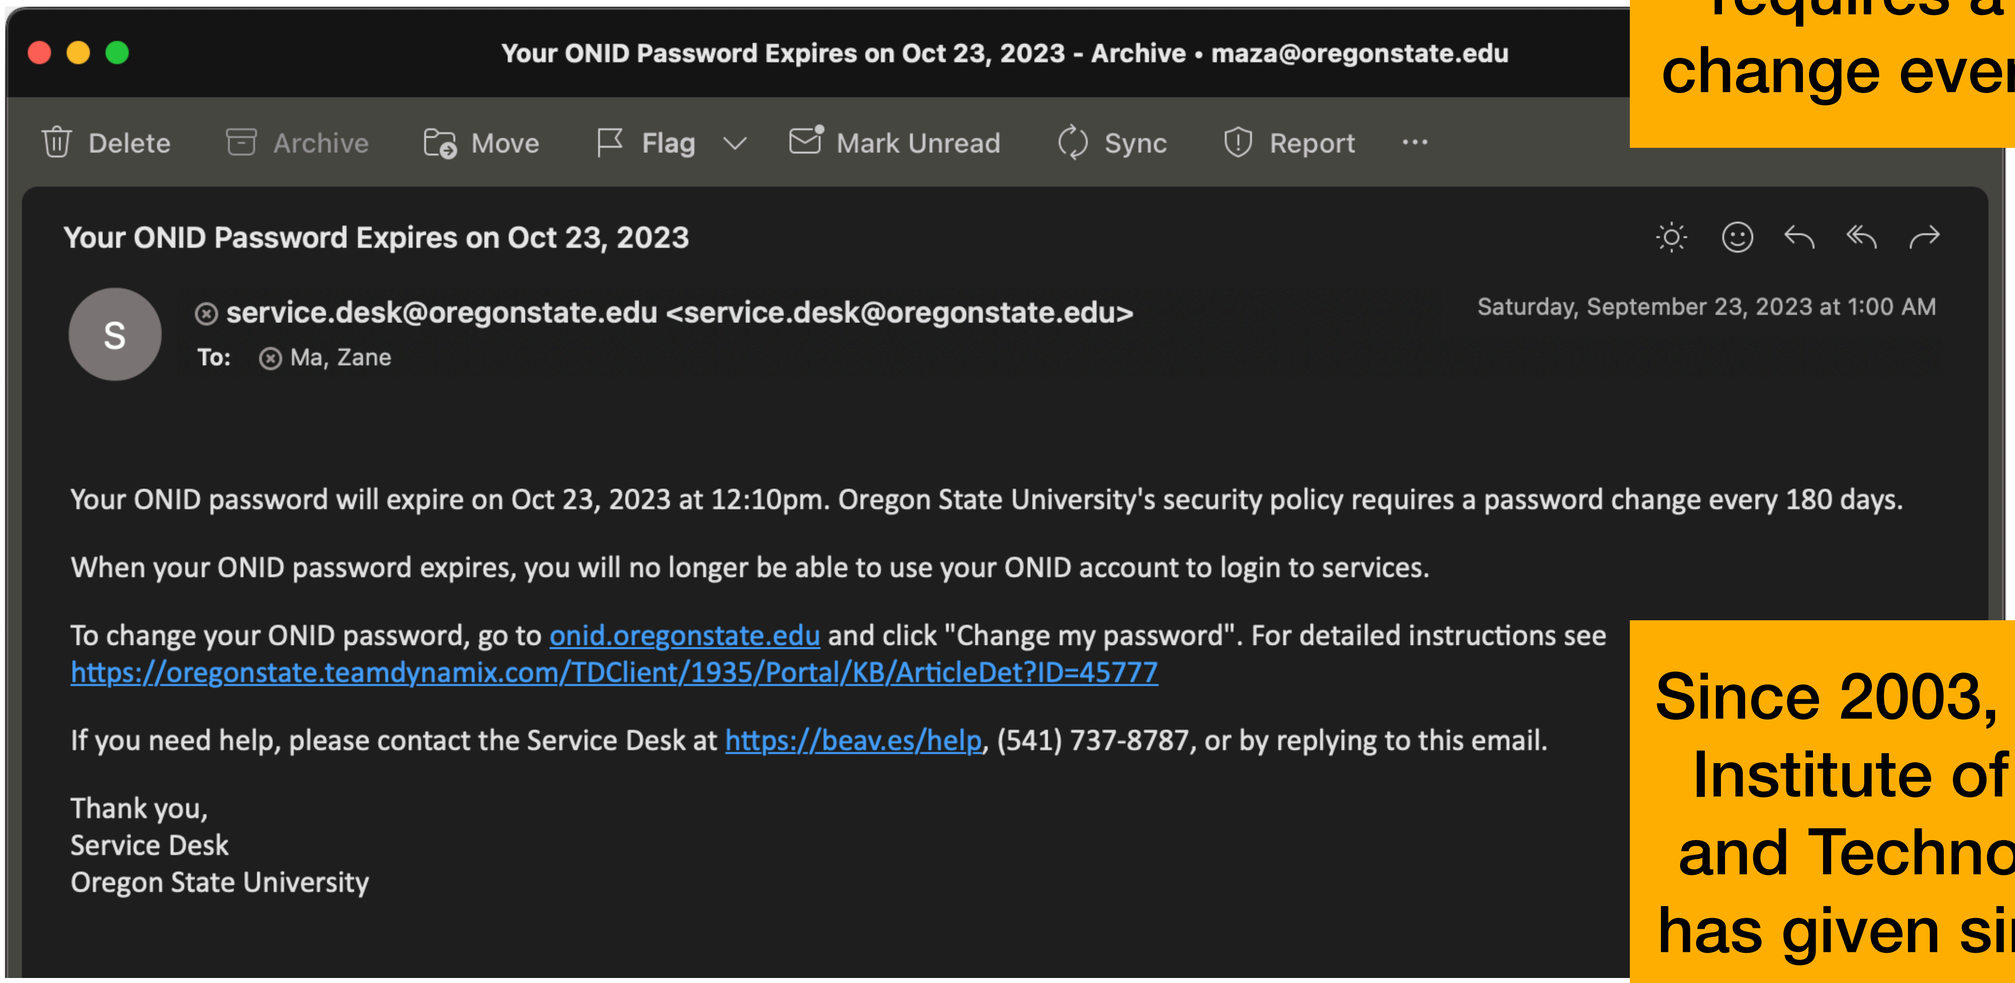

Since 2003, US National Institute of Standards and Technology (NIST) has given similar advice

Why periodically expire passwords?

"Changing passwords frequently narrows the window within which an account is usable to an attacker before he has to take additional steps to maintain access."

"Password expiration does not offer any benefit when an attacker wants to do all of the damage that he's going to do right now. It does offer a benefit when the attacker intends to continue accessing a system for an extended period of time."

S. Alexander, Jr. In defense of password expiration. Post to League of Professional System Administrators (LOPSA) blog, April 2006.

UNC Chapel Hill in 2010

Similar password policy for their single-sign-on system (i.e., ONYEN)

- Required to change password **every 3 months**
- Password cannot have been used for the account in the last year
- Password must be at least 8 characters long, contain >= 1 letter + 1 digit
- Password must contain >= 1 special character
- Password must share < 6 consecutive characters with the username
- Password must not start with a hyphen, end with a backslash, start or end with a space, or contain a double-quote anywhere except as the last character

Zhang, Yinqian, Fabian Monrose, and Michael K. Reiter. "The security of modern password expiration: An algorithmic framework and empirical analysis." CCS. 2010.

Dataset

51K unsalted MD5 password hashes for 10K defunct ONYEN accounts (2004-2009)

Password hash? A deterministic, one-way transformation: hash(password) = <random-looking password hash>. Since the hash function is basically impossible* to reverse, only storing hashed passwords is a way to slow down adversaries that compromise a password database

Salting? Pre-computed “rainbow tables” (hashes of popular passwords) or brute-force cracking can be somewhat effective. To hinder adversaries even more, we store <hash(salt + password), salt> instead of <hash(password)>

Dataset

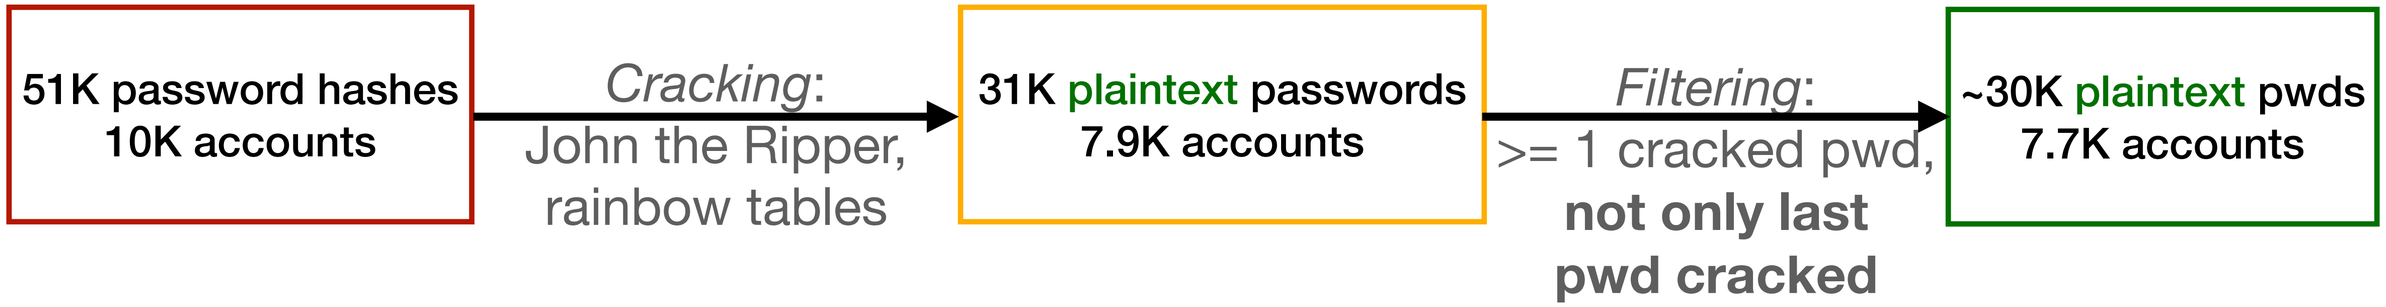

John the Ripper = software that brute-force calculates hashes for passwords

Rainbow table = lookup table for known hash —> password

Why this filtering requirement? Any others?

Cracked Passwords

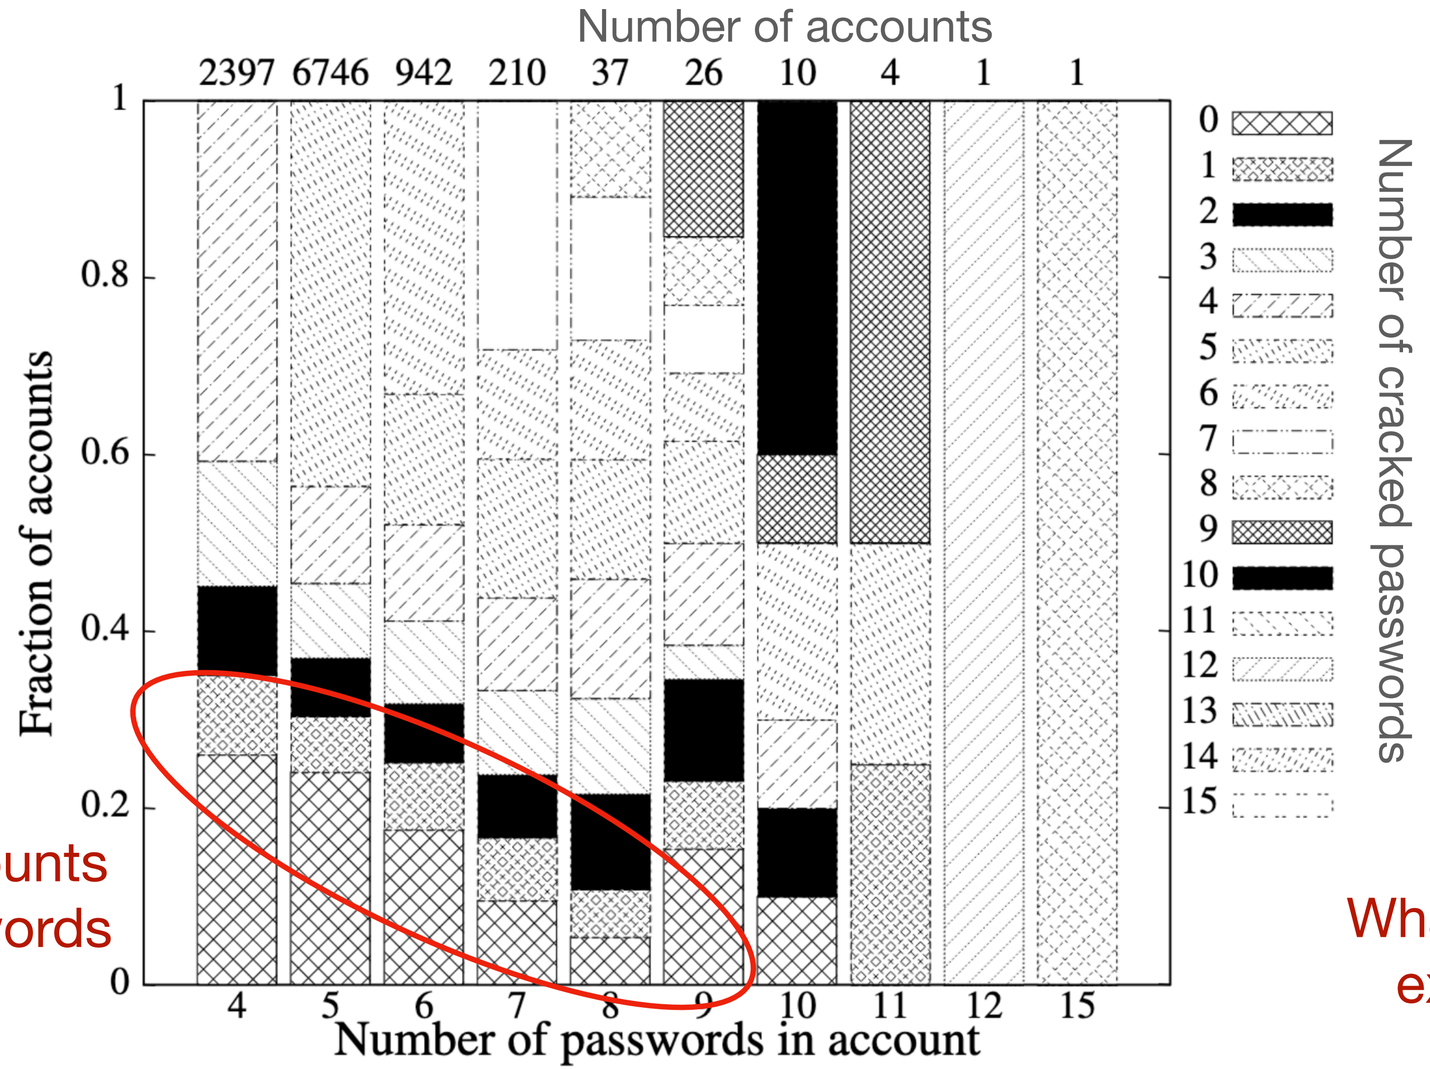

Decreasing % of accounts with uncracked passwords

What are possible explanations?

Threat Model

- Meta: the assumptions of capabilities and motives for both attacker + defender

- Threat Model 1: The goal of password expiration is to protect against “an attacker who has acquired a valid password” - how much does changing passwords deter such an attacker?

- Threat Model 2: For an arbitrary attacker who doesn’t necessarily already have access to a valid password, does changing passwords make it easier / harder to guess passwords?

Experiment for Threat Model 1

- Threat Model 1: an attacker who has already acquired an old password
- Hypothesis: based on prior qualitative studies, users reported that new passwords are often related to the prior password; thus, they can be easily guessed
- First experiment/analysis: look at string-similarity metrics for ONYEN accounts
- Next-level experiment: try to crack passwords without knowledge of old password, compare success rates with cracking attempts that **do** know old password

Password prediction: Transform tree

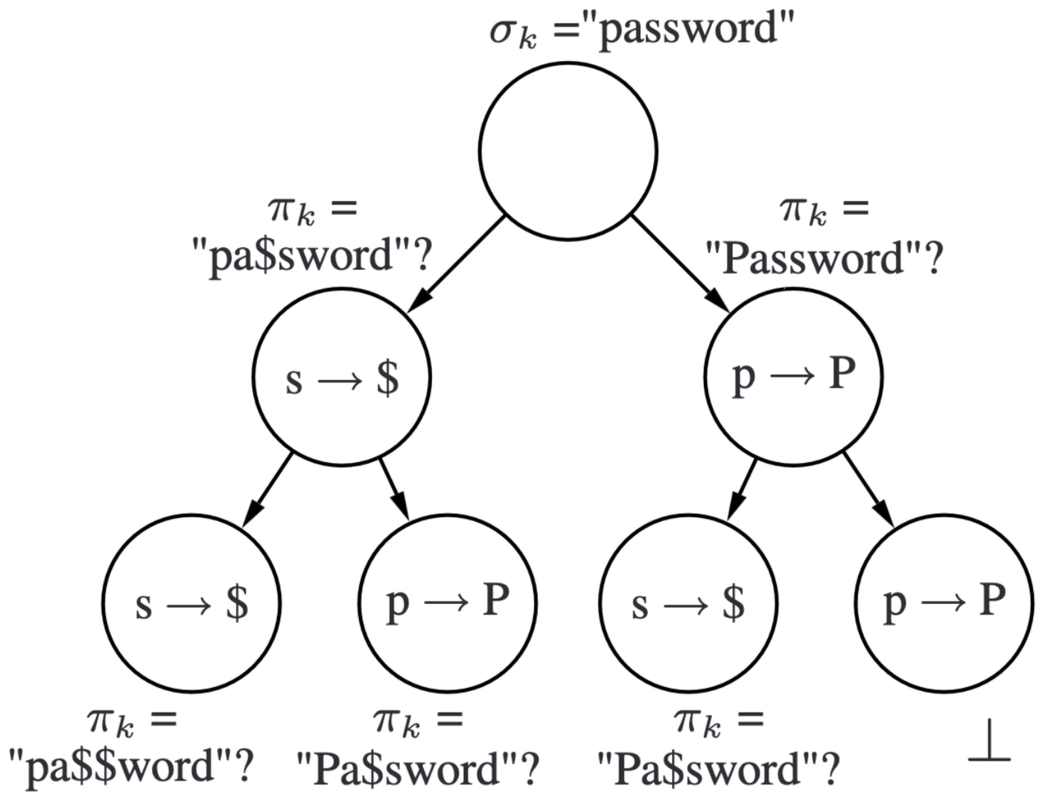

- No prior cracking algorithm based on specific password(s) for an account
- Model each transformation as a node in a tree, which yields a modified password π_k or error $\perp$
- High-level insight: use data from many accounts to identify popular transformations, and try to guess new passwords based on known old passwords

Results for Threat Model 1

- Further subdivide the threat model into two groups
- Threat Model 1A: offline attacks (e.g., password-encrypted file) that have unlimited password guess attempts
    - ~41% of passwords can be broken from an old password in < 3 seconds
- Threat Model 1B: online attacks (e.g., website login) that have limited guess attempts
    - ~17% of accounts can be broken in under five online guesses

Research Recap

- Collected + cracked + cleaned up dataset of account password hashes
- Threat model: attacker with knowledge of a valid old password
- New algorithm: using transform trees to predict passwords based on prior
- Results
    - 41% of new password can be guessed from old password in 3 seconds
    - 17% of new passwords can be guessed in 5 attempts

Wishlist

- What else do we want to know?

- Ways to improve on the transform tree; new transform types, better learning algorithm, bigger/better data
- Explicit comparison of Threat Model 1 with / without password expiration
- Threat Model 2! Attacker w/o old password - do expiration policies make passwords less secure in general?
- Should we make password change policies stricter (e.g., new password can't be similar to old password) or should we eliminate password expiration?

Back to password expiration…

OSU's security policy requires a password change every 180 days.

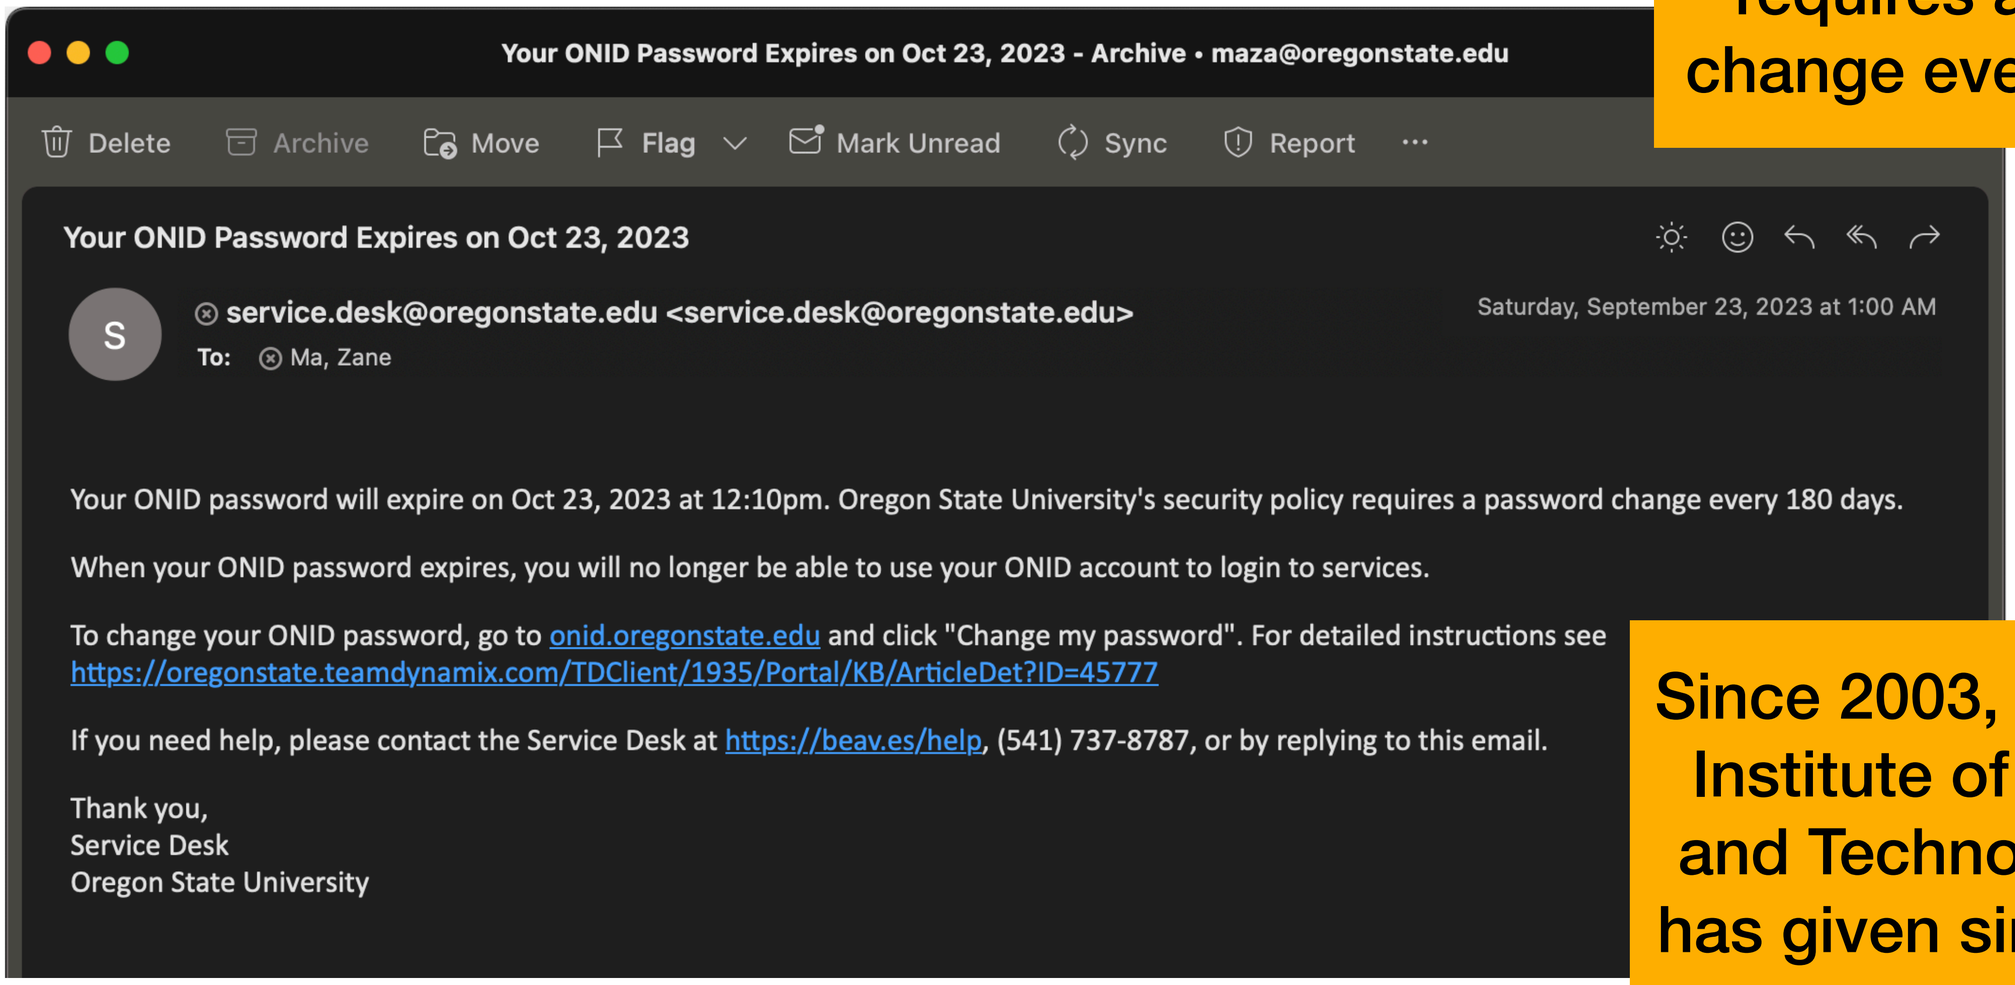

Since 2003, US National Institute of Standards and Technology (NIST) has given similar advice

Challenging theories with empiricism

- Existing policy / theory: changing passwords is good, like changing door locks
- UNC study showed: changing passwords doesn't protect well against attacker with access to old passwords
- Carleton study showed: Users who know they will have to change their password do not choose strong passwords to begin with and are more likely to write their passwords down
- CMU study: CMU students, faculty and staff who reported annoyance with the CMU password policy ended up choosing weaker passwords than those who did not report annoyance

NIST changed their recommendations in 2016
OSU still has 180-day password changes

Taking a step back…

- What is this term “empiricism”? What does it really mean?

To answer this question, we must start with:

- What is Science?
- How does empiricism contribute to Science?
- Is Computer Security scientific? Can we make it scientific?
- What makes Science of Security hard?

SoK: Science, Security and the Elusive Goal of Security as Scientific Pursuit

Cormac Herley
Microsoft Research

P.C. van Oorschot
Carleton University

2017 IEEE S&P

What do we mean by “science”?

- Equations?
- Numbers / Graphs?
- Repeatable Experiments?
- Rigor? Proofs?
- Scientific method?


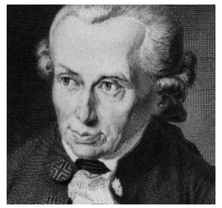
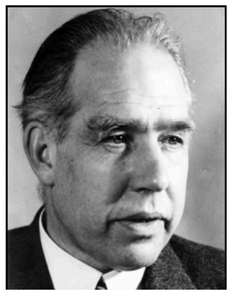
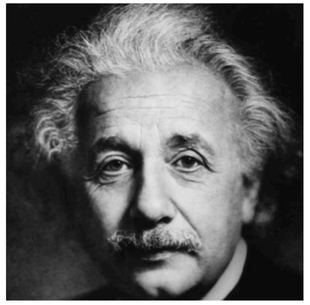
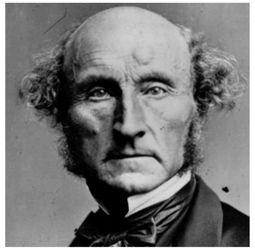
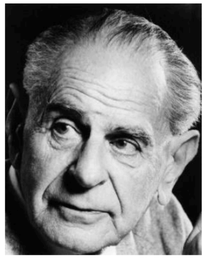
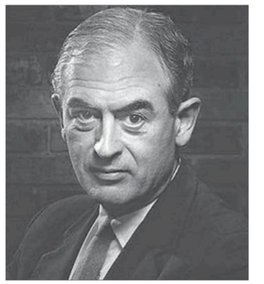
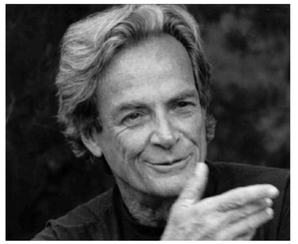

Why do we want science?
What are the desirable properties?

What's the consensus from others?

If theory conflicts with observation, it's wrong.

Conflict with observation must be possible, implying:

1. Science is induction, not deduction
2. Claims must be falsifiable

Induction vs Deduction

Induction: statements about real world (always uncertain) based on observation

Deduction: proved-true statements from axioms

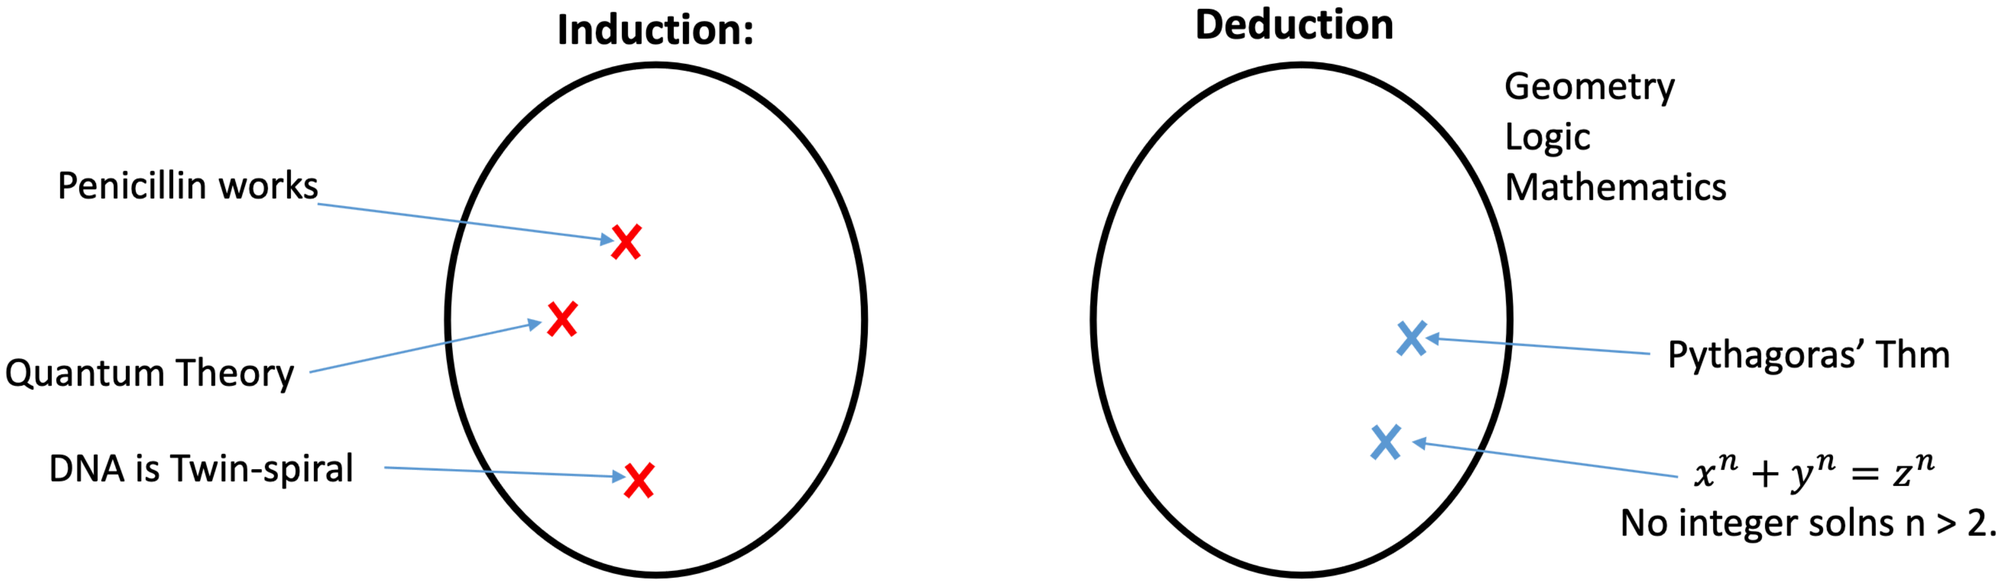

	Inductive Statements	Deductive Statements
Describe real-world?	Yes	No

Induction vs Deduction

Induction: statements about real world (always uncertain) based on observation

Deduction: also, the application of logic to inductive claims/assumptions

Induction

Speed of object falling in a vacuum = g * t

Deduction

Rate of speed change = acceleration = g

Why do we believe inductive assumption?
What makes it scientific?

Falsifiability

"A theory which is not refutable by any conceivable event is non-scientific. Irrefutability is not a virtue of a theory (as people often think) but a vice." - K. Popper

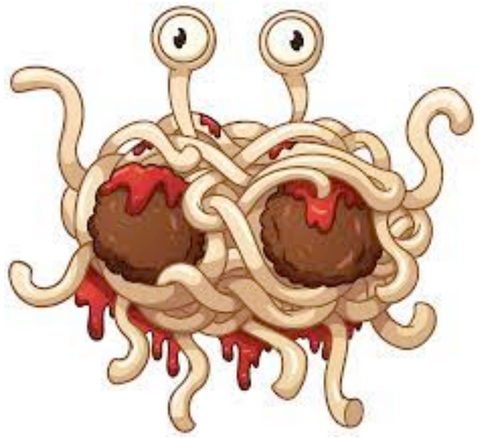

If X cannot be falsified by any observation then:

1. X is consistent with every possible observation
2. Nothing observable (i.e., the real-world) depends on X

	Inductive Statements	Deductive Statements
Describe real-world?	Yes	No
Believe when:	Try to falsify and fail	Have a proof

Wait, Math isn't science?!

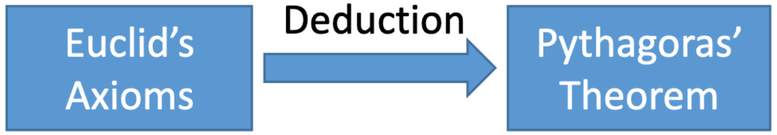

No observation contradicts Pythagoras' Theorem.

- If $a^2 + b^2 \neq c^2$ for the door we don't say theorem wrong.

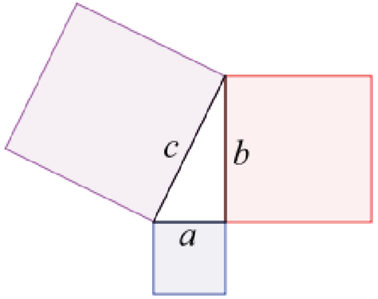

Axiom: parallel lines meet at infinity

Assumption: attacker can't do log in a finite field

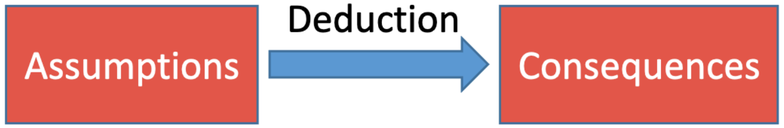

Observations contradicting assumption are possible. Scientific claims retain uncertainty

Whether a real-world system satisfies assumptions is an empirical claim (and must be tested).

How does this relate to computer security?

1. Failure to separate induction/deduction

Example Problematic Claim: *"There is no (and cannot be) empirical evidence for the security of a design. […] The only way to do so is to develop a formal mathematical model and language in which to reason about such schemes."*

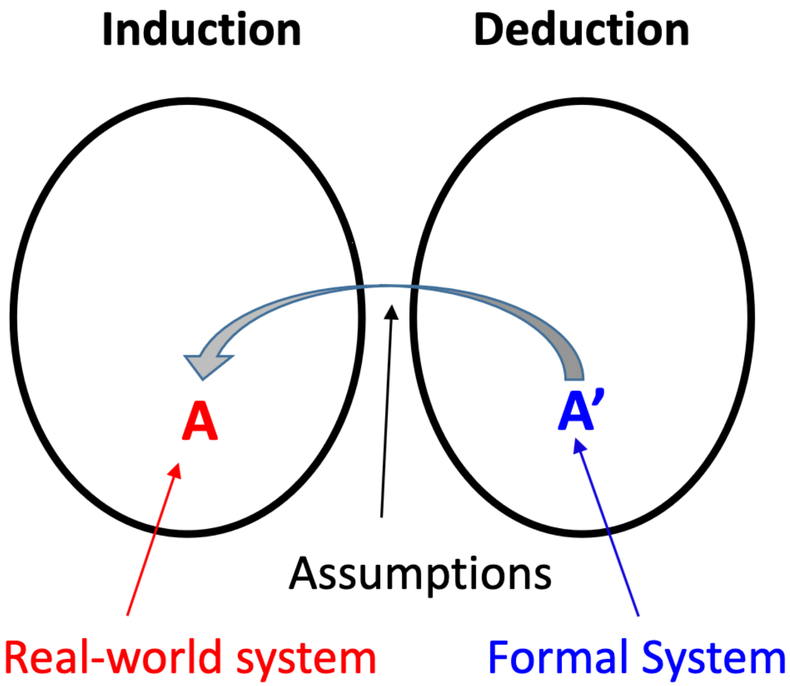

Formal System A':

- Proof

Real-world System A:

- Proof + argument that assumptions match reality

Whether assumptions match reality can only be explored empirically

1. Failure to separate induction/deduction

Example:

- **A':** attack on SSL must solve hard problem
- **A:** Remote Timing Attacks are Practical (2003)

- **A** enjoys properties of **A'** is assumed, not proved
- No possibility of proving **A** immune to attack
- No end-run around messiness of real-world

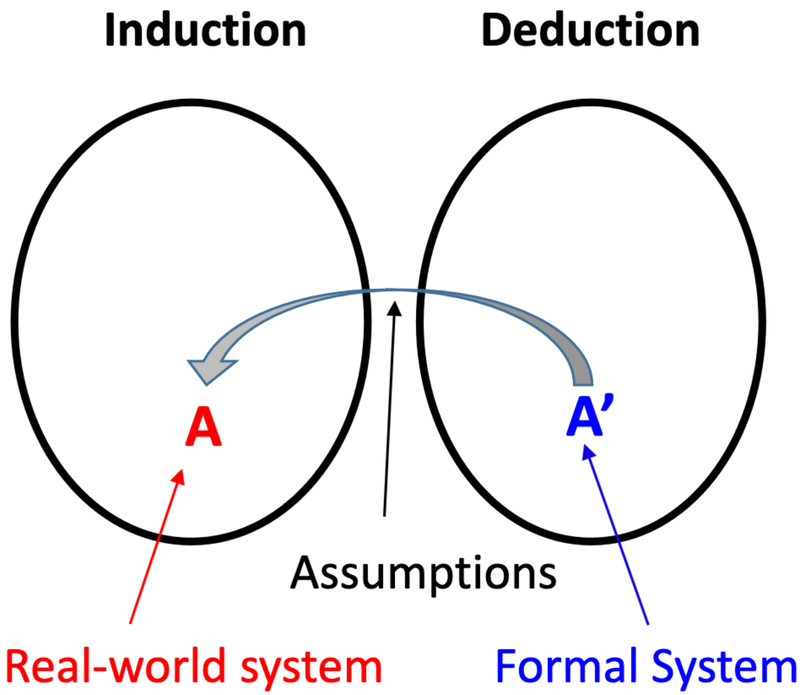

A proof + argument that assumptions are reasonable is not a proof.
It's also not scientific w/o attempts to refute assumptions.

2. Failure to challenge theory with observation

"Passwords should contain a mix of upper, lower and special chars."

Morris&Thompson, 1979

Three+ decades of ***assuming*** this leads to more guess-resistant pwds.

What's the basis for claiming:

- *Passwords should be changed every 90 days.*
- *Should always obey browser warnings*

Do we have A/B tests? Observations of improved outcomes?

3. Reliance on implicit / unfalsifiable assumptions

When making claims such as:

- Changing passwords every 90 days improves outcomes
- Choosing stronger password improves outcomes

We should know what specific evidence would refute these statements!

3. Reliance on implicit / unfalsifiable assumptions

When making claims such as:

- Changing passwords every 90 days improves outcomes
- Choosing stronger password improves outcomes

We should know what specific evidence would refute these statements!

Claim: System X is secure / insecure

Is this scientific? If not, can we make it scientific?

Is science a reasonable goal for security?

Claims that unique aspects of security exempt it from a scientific approach are unhelpful.

- Common excuses: “But active adversary, no fundamental laws, man-made artifacts…..”
- Science is the best way we know of making inferences in the real world
    - Acknowledgment of fallibility —> self-correction

What makes Science of Security hard?

- Security often deals with *sensitive* data; tricky ethical access to data
- Security is a process that evolves over time; people change, hardware changes, threats / defenses change
- Security is a broad field across all areas of computing - finding foundational security principles requires lots of empirical evidence + deep analysis / insight

- Science is hard! Requires knowledge of how to measure, big data, statistics, in addition to deep domain knowledge (to understand assumptions / how to falsifiability)

How empiricism contributes to research

- Challenges existing assumptions about security (e.g., 1024-bit RSA is enough)
- Uncover new implicit assumptions (e.g., Mining p’s and q’s paper)
- Identify new theories from measurement (e.g., economics of security)

TODOs for you

Specify presentation preferences by **Wednesday, October 4th**. Sign-up link on the syllabus at https://empirical-security.net/syllabus

Student-presentations begin **October 11th** - I will reach out to students to schedule a time to meet

Create a project team by **Friday, October 6th**. See Canvas discussion thread